

Möbius Transformations For Global Intrinsic Symmetry Analysis

Vladimir G. Kim, Yaron Lipman, Xiaobai Chen, and Thomas Funkhouser

Princeton University

Abstract

The goal of our work is to develop an algorithm for automatic and robust detection of global intrinsic symmetries in 3D surface meshes. Our approach is based on two core observations. First, symmetry invariant point sets can be detected robustly using critical points of the Average Geodesic Distance (AGD) function. Second, intrinsic symmetries are self-isometries of surfaces and as such are contained in the low dimensional group of Möbius transformations. Based on these observations, we propose an algorithm that: 1) generates a set of symmetric points by detecting critical points of the AGD function, 2) enumerates small subsets of those feature points to generate candidate Möbius transformations, and 3) selects among those candidate Möbius transformations the one(s) that best map the surface onto itself. The main advantages of this algorithm stem from the stability of the AGD in predicting potential symmetric point features and the low dimensionality of the Möbius group for enumerating potential self-mappings. During experiments with a benchmark set of meshes augmented with human-specified symmetric correspondences, we find that the algorithm is able to find intrinsic symmetries for a wide variety of object types with moderate deviations from perfect symmetry.

1. Introduction

Detecting symmetries in 3D surface meshes is an important research problem with applications in object alignment, matching, compression, completion, and beautification.

Several effective methods have been proposed for detecting extrinsic symmetries (e.g., [MGP06, PSG*06]). They generally utilize a low-dimensional representation for extrinsic symmetry transformations (e.g., rigid transformations have six degrees of freedom) to perform a search for the best symmetry transformation using variants of the Generalized Hough transform, Geometric Hashing, RANSAC, ICP, or some other rigid alignment algorithm.

Due to the difficulty in parameterizing non-rigid transformations, there has been fewer advances in detecting intrinsic symmetries – ones that are invariant to isometric deformations. Recently, methods have been proposed to find symmetric feature correspondences with embeddings based on eigenvectors of the Laplace-Beltrami operator [OSG08], loopy Bayesian belief propagation [LTSW09], subgraph matching [BBW*08], and combinatorial search of feature correspondences [RBB*10]. However, these methods are either theoretically intractable [RBB*10], expensive in practice [LTSW09], depend on local shape similarities [BBW*09a], and/or fail to detect all types of intrinsic symmetries [OSG08].

The focus of our work is to investigate algorithms for

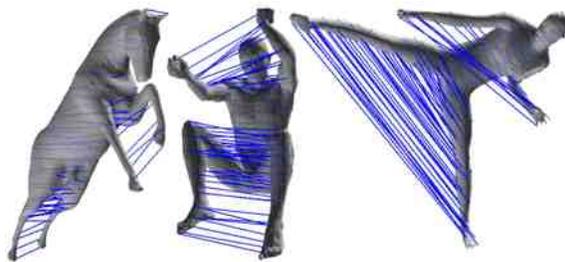


Figure 1: Intrinsically symmetric point correspondences found by our algorithm.

discovery of intrinsically symmetric point correspondences using methods based on analysis of Möbius transformations [LF09, ZWW*10] (Figure 1). The Möbius transformation group can describe all isometric transformations with a low dimensional representation (six degrees of freedom for genus-zero surface), and thus it can be searched efficiently for a global mapping from a surface onto itself using algorithms traditionally used for finding extrinsic symmetries.

In this paper, we explore the idea of using Möbius transformations to find correspondences between intrinsically symmetric points on a 3D mesh. In addition to the main idea, we make several research contributions: 1) we develop theory regarding the relationship between intrinsic symmetry

and symmetry-invariant functions, 2) we utilize that theory to extract point sets from 3D meshes that provably contain stationary points and correspondences with respect to an intrinsic symmetry transformation if one exists, 3) we observe that intrinsic symmetries in surfaces are contained in the larger group of anti-Möbius transformations, 4) we search triplets and quadruplets of point correspondences to find the anti-Möbius transformation that best maps a mesh onto itself, 5) we utilize properties of intrinsic symmetry to cull anti-Möbius transformations that are obviously not self-isometries, and 6) we evaluate the anti-Möbius transformations found with our algorithm with respect to human-annotated symmetric point correspondences for three benchmark data sets containing 366 meshes in total.

2. Related work

Understanding the symmetries of shapes is a well studied problem in many disciplines, including perceptual psychology, computer vision, computational geometry, and computer graphics.

Most prior work has focused on detection of extrinsic symmetry transformations – rigid body transformations that map a surface onto itself. For example, methods have recently been proposed to detect partial, extrinsic symmetries using the Generalized Hough Transform [MGP06], Geometric Hashing [GCO06], and spectral analysis of an association graph [LE06]. Generally speaking, they sample the space of possible transformations, measuring how well or how much of the surface is aligned for each transformation or correspondence. They are effective and efficient because the space of transformations has low dimensionality and thus can be sampled efficiently (e.g., planar reflections have only three degrees of freedom).

More recently, work has focused on detecting intrinsic symmetry transformations – i.e., isometric transformations that map a surface onto itself. Due to the difficulty in parameterizing non-rigid transformations, most recent work on this problem has focused on combinatorial algorithms that search for symmetric feature correspondences [BBW*08, BBW*09b, LTSW09, RBB*10, WJG07, ZSCO*08]. This problem is NP-Hard in theory, but several approaches are practical when shape descriptors can be used to prune potential feature correspondences.

For example, in the work of Raviv et al [RBB*10], feature points are generated by iteratively selecting the point farthest from previously selected samples. Then, a histogram is computed for each feature point that stores the distribution of geodesic distances (or diffusion distances [RBB*10]) to other feature points, and a branch-and-bound algorithm is used to search for sets of feature point correspondences where the histograms for corresponding points are compatible. The resulting sparse set of point correspondences is refined with the GMDS algorithm [BBK08], possibly augmented with a fuzzy set membership function [RBBK10].

To avoid searching the exponential complexity of the correspondence space, the method uses a threshold on the similarity between histograms of geodesic distances to establish a sparse set of point correspondences. As a result, it is only practical for cases where the self-mapping is not too far from isometric.

An alternative approach is to embed a surface into a high-dimensional Euclidean space where intrinsic symmetries appear as extrinsic ones. For example, Ovsjanikov et al. perform an embedding into a signature space defined by the eigenfunctions of the Laplace-Beltrami operator [OSG08]. This embedding transforms intrinsic symmetries into extrinsic ones. However, since the embedding space is high-dimensional, it is difficult to detect extrinsic symmetry transformations in that space directly. Instead, Ovsjanikov et al. consider a reduced embedding space, where each point is represented only by a sequence of signs of non-repeating eigenfunctions. As a result, their approach is sensitive to shape perturbations that flip or re-order eigenspaces.

Some researchers have proposed algorithms for detecting structures that characterize intrinsic symmetries, without explicitly establishing a set of symmetric point correspondences. For example, Xu et al. proposed an algorithm to detect the Partial Intrinsic Reflectional Symmetry Axis of a surface, the locus of points that are stationary with respect to a reflectional symmetry [XZT*09]. Lipman et al. proposed a method for detecting symmetry orbits, sets of points that map onto one another within each symmetry group [LCDF10]. These methods solve a subset of the problem addressed in this paper – our method finds all symmetric point correspondences, which includes stationary points and orbits.

Our algorithm is motivated by the Möbius Voting paper of Lipman and Funkhouser [LF09]. In that paper the authors suggest searching for (near-)isometries between two different genus-zero surfaces \mathcal{M}, \mathcal{N} in the space of Möbius transformations, which contains the entire collection of (orientation preserving) conformal maps between the surfaces. They leverage the low dimensionality of Möbius transformations to find point correspondence between two surfaces with the following algorithm. First, both meshes are mapped to the extended plane $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ with the mid-edge flattening scheme. Second, the meshes are sampled with the farthest point algorithm to produce uniform sample set on each mesh, $\mathcal{S}_{\mathcal{M}} \subset \mathcal{M}, \mathcal{S}_{\mathcal{N}} \subset \mathcal{N}$. Third, triplets of potential point correspondences are generated randomly, $z_i \in \mathcal{S}_{\mathcal{M}}, w_i \in \mathcal{S}_{\mathcal{N}}, i = 1, 2, 3$. For each triplet, the unique Möbius transformation $m(\cdot)$ interpolating the correspondences $m(z_i) = w_i, i = 1, 2, 3$ is used to map $\mathcal{S}_{\mathcal{M}}$ into $m(\mathcal{S}_{\mathcal{M}})$. Mutually closest points are extracted (in the flattened space). If the fraction of mutually closest is above a threshold, a deformation error (the reciprocal of the average Euclidean distance between mutually closest pairs) is computed and summed into a fuzzy correspondence matrix. After enough votes have been cast, correspondences are extracted from the fuzzy cor-

respondence matrix using a greedy algorithm. Finally, good correspondences with low confidence (values in correspondence matrix) pairs are further extracted using geodesic consistency.

Our work is different from Möbius Voting in four important ways. First, we propose novel theory and algorithms for extracting stable intrinsically symmetric sets of points from a surface (to be used as candidates for correspondence). Second, we observe that intrinsic symmetries in surfaces are contained in the larger group of anti-Möbius transformations (anti-Möbius transformations are both “regular” Möbius transformation and Möbius transformations composed with a reflection) and then adjust the algorithm to consider this larger set. Third, we introduce methods for pruning potential triplets of candidate correspondences that take advantage of stationary points and other properties unique to symmetry. Finally, we improve the implementation for almost every step of the original algorithm.

3. Method

The goal of our work is to detect intrinsic symmetries in surface meshes. Intrinsic symmetries are self-mappings of smooth surface \mathcal{M} to itself $f : \mathcal{M} \rightarrow \mathcal{M}$ which preserve pairwise geodesic distances $d_g(p, q) = d_g(f(p), f(q))$. In other words they are self-isometries of the surface and preserve the metric tensor.

In this work, we assume \mathcal{M} is orientable and topologically a sphere (genus zero). We note however that the following arguments can be generalized to arbitrary orientable surfaces [GY03], but this is outside the scope of this paper. Isometries $f : \mathcal{M} \rightarrow \mathcal{M}$ can be either globally orientation preserving or globally orientation reversing. Intuitively, orientation preserving means that f preserves the normal direction, while orientation reversing means that f flips the normal direction. An orientation preserving isometry can be viewed as a particular instance of a conformal map, while orientation inverting isometry is an instance of anti-conformal map. This means that for finding (perfect) symmetries $f : \mathcal{M} \rightarrow \mathcal{M}$ it is enough to search within the collection of conformal and anti-conformal mappings $m : \mathcal{M} \rightarrow \mathcal{M}$.

Motivated by these observations we have developed an algorithm to detect intrinsic symmetry in meshes. Our algorithm proceeds in the following three steps:

1. **Symmetric point sampling.** We extract two symmetry-invariant sets of points on the surface $\mathcal{S}_1 \subset \mathcal{S}_2 \subset \mathcal{M}$, where \mathcal{S}_1 is “coarse” and \mathcal{S}_2 is “dense”. By symmetry-invariant set $\mathcal{S}_i, i = 1, 2$ we mean that for all symmetries $f : \mathcal{M} \rightarrow \mathcal{M}$ and all $p \in \mathcal{S}_i, f(p) \in \mathcal{S}_i$. Since f^{-1} is also symmetry this implies $f(\mathcal{S}_i) = \mathcal{S}_i$. The coarse set \mathcal{S}_1 is used to produce candidate anti-Möbius transformations, while \mathcal{S}_2 to assign correspondences.
2. **Searching anti-Möbius transformations.** The next stage iteratively picks small subsets of points (generating

sets) from the coarse set \mathcal{S}_1 , where each generating set provides constraints to define an anti-Möbius transformation that is potentially an intrinsic symmetry. Each candidate transformation is assigned with an alignment score based on how well the transformed surface aligns with itself. After exhausting all possible generating sets from \mathcal{S}_1 the algorithm returns the best aligning anti-Möbius transformation.

3. **Extracting correspondences.** Using the best anti-Möbius transformation from previous step we define a set of correspondences within the set \mathcal{S}_2 by taking mutually closest pairs. The last step is to use the extracted correspondences within set \mathcal{S}_2 to predict more correspondences within the set \mathcal{S}_2 . The prediction is made based on consistency of geodesic distances to found correspondences.

In the following sections we provide the details for each of these steps.

4. Symmetric point sampling

The goal of the first step is to construct an algorithm that extracts symmetry-invariant point samples on the surface. A symmetry invariant point set $\mathcal{S} \subset \mathcal{M}$ satisfies $f(\mathcal{S}) = \mathcal{S}$, for all global intrinsic symmetries $f : \mathcal{M} \rightarrow \mathcal{M}$. Our algorithm uses sets with this property to form candidate symmetric point correspondences in the second and third steps.

The key idea in constructing the symmetry-invariant set $\mathcal{S} = \{p_i\} \subset \mathcal{M}$ is to extract critical points of a symmetry invariant function $\mathcal{S} = \{p \in \mathcal{M} \mid \nabla_p \Phi = 0\}$, where symmetry-invariant function $\Phi : \mathcal{M} \rightarrow \mathbb{R}$ is defined by $\Phi(f(p)) = \Phi(p)$, for all symmetries f and $p \in \mathcal{M}$. Intuitively, the collection of critical points of symmetry-invariant function is a symmetry-invariant set of points. This is proved in Section 7. In that section we also prove more properties of this set and its relation to intrinsic symmetries.

Using this principle, we construct a two stage algorithm for sampling symmetry-invariant point sets on surfaces. First, we construct a coarse set of symmetry-invariant points, and second, we add points to this set while maintaining the invariance property.

Coarse symmetry-invariant set. We would like to take a symmetry-invariant function $\Phi : \mathcal{M} \rightarrow \mathbb{R}$ and extract its critical points. Since we want our algorithm to be robust to deviations from perfect symmetries and non-symmetric extraneous and missing parts, it is desired to use symmetry-invariant functions which involve averaging over many geodesic distances. A natural function in this context is the Average Geodesic Distance (AGD):

$$\Phi_{agd}(p) = \int_{\mathcal{M}} d_g(p, q) d\text{vol}_{\mathcal{M}}(q). \quad (1)$$

Proposition 7.2 in Section 7 provides a proof that AGD is indeed symmetry-invariant function. The coarse symmetry-

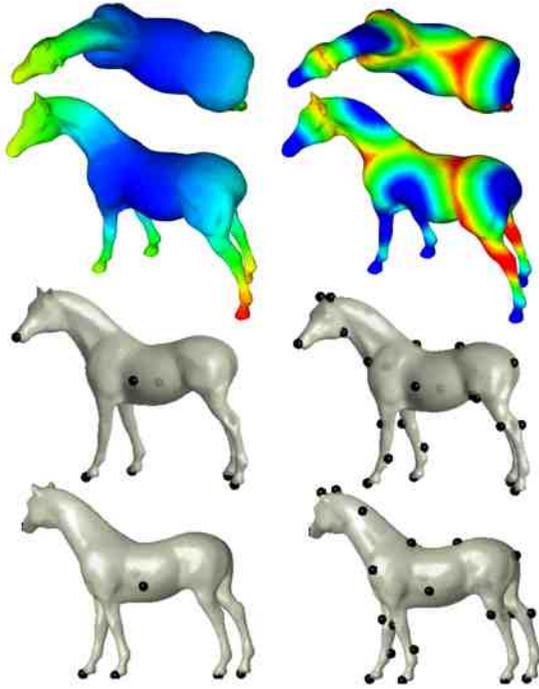


Figure 2: The symmetric sampling procedure is demonstrated on an intrinsically bilateral reflective horse model. In the left column we show the AGD function (top) and its extracted critical points \mathcal{S}_1 (bottom, two opposite views). In the right column the MGD to \mathcal{S}_1 is shown on top, and the augmented symmetry-invariant set \mathcal{S}_2 is shown on bottom (two opposite views). Note the symmetric sampling.

invariant set is then defined

$$\mathcal{S}_1 = \{p \in \mathcal{M} \mid \nabla |p \Phi_{agd} = 0\}. \quad (2)$$

Figure 2 (left column) shows the AGD and the corresponding \mathcal{S}_1 for an intrinsically symmetric horse (bilateral reflective symmetry).

Augmenting the symmetry-invariant set. Once a coarse symmetry-invariant set \mathcal{S}_1 is found, we can use it to find a superset $\mathcal{S}_2 \supset \mathcal{S}_1$ that is also symmetry-invariant. For that end we will construct a second symmetry-invariant function based on the Minimal Geodesic Distance (MGD) to the symmetry-invariant set \mathcal{S}_1 :

$$\Phi_{mgd}(p) = \Phi_{mgd}(p; \mathcal{S}_1) = \min_{q \in \mathcal{S}_1} d_g(p, q). \quad (3)$$

Proposition 7.3 in Section 7 provides a proof of the invariant property of this function. We define $\mathcal{S}_2 = \{p \in \mathcal{M} \mid \nabla |p \Phi_{mgd} = 0\}$. This procedure can be repeated by extracting critical points of $\Phi_{mgd}(p; \mathcal{S}_2)$ and defining this set as an updated \mathcal{S}_2 . This way we can

produce arbitrary large set \mathcal{S}_2 . In our implementation, for stability, we choose only local maximas of $\Phi_{mgd}(p; \mathcal{S}_1)$, which are the vertices of the Voronoi diagram defined by the centers \mathcal{S}_1 . Figure 2 (right column) shows the MGD function for \mathcal{S}_1 and its augmented set \mathcal{S}_2 (after 2 iterations). The set \mathcal{S}_2 is not guaranteed to be uniformly spread across the surface, since it depends on distances to points in the coarse set \mathcal{S}_1 . To increase uniformity a possible modification of this algorithm is to take only critical points not far from the global maximum: $\mathcal{S}_2 = \{p \in \mathcal{M} \mid \nabla |p \Phi_{mgd} = 0 \text{ and } \Phi_{mgd}(p) > \tau \cdot \max_{\mathcal{M}} \Phi_{mgd}\}$, with $\tau \in [0, 1]$.

Implementation on meshes For calculating the symmetry invariant sets on meshes $\mathcal{M} = (V, E)$, where V is the vertices set and E the edges set, we have discretized the above concepts in the standard way. First, we approximate the geodesic distances on surfaces with Dijkstra's algorithm. Second, the average geodesic distance function is approximated by the average distance to all vertices, that is

$$\Phi_{agd}(v_i) = \sum_{v_j \in V} \frac{A_i}{3} d_g(v_i, v_j),$$

where A_i is the area of the 1-ring of vertex i . In our implementation for extracting the critical points we have smoothed the AGD function by averaging it over local 1-rings, and picked local (1-ring again) extremas (maxima and minima).

5. Searching Möbius transformations

The second step is to search for the anti-Möbius transformation that induces the best mapping of the mesh onto itself. The anti-Möbius group is the group of Möbius transformations augmented by Möbius transformations composed with a reflection ($z \mapsto \bar{z}$, as explained below in more details).

Previous work [LF09] has suggested to search the Möbius transformations group for finding isometries and near-isometries between two surfaces. Our work builds upon this idea once making the observations that intrinsic symmetries are self-isometries of the surface. We suggest a few modifications over the algorithm suggested in [LF09] mainly to take advantage of the fact that we are looking for symmetries in one surface and to provide more justified choices for some components of the algorithm.

The outline of our algorithm for searching the anti-Möbius transformations is as follows:

1. **Uniformization.** We map the mesh surface to the extended complex plane $\hat{\mathbb{C}}$ via the mid-edge uniformization procedure, as in [LF09].
2. **Generating sets of anti-Möbius transformation.** We consider correspondences amongst triplets and quadruplets of points from the symmetry-invariant set \mathcal{S}_1 to generate candidate anti-Möbius transformations.
3. **Pruning generating sets.** For each generating set, we apply a set of filters to determine whether the distances

and properties of corresponding points are conservatively consistent with a symmetric arrangement.

4. **Measuring alignment score.** If the generating set is not pruned, we apply the implied anti-Möbius transformation to a dense set of points S_u and measure how well they align to estimate an alignment score.
5. **Return best anti-Möbius transformation.** As generating sets are considered, we remember the anti-Möbius transformation that yielded the (few) best alignment score, and return it as the final answer.

Uniformization. Uniformization theory guarantees that (smooth enough) genus zero surfaces can be mapped conformally onto the extended complex plane $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ which is conformally equivalent to the 2-sphere $S^2 = \{(x, y, z) \mid x^2 + y^2 + z^2 = 1\} \subset \mathbb{R}^3$. (The case of higher genus surfaces can be treated using covering spaces [GY03], but we did not treat this case in our algorithm.)

Uniformization or conformal parameterization of surfaces is an active research field and several researchers have suggested different techniques. In this work we chose to use the mid-edge uniformization technique suggested in [PP93, LF09] due to its simplicity. Employing their technique results in embedding of each mid-edge point in the complex plane. Then, we map each vertex $v_i \in V \mapsto z_i \in \hat{\mathbb{C}}$ to the average of its 1-ring mid-edge points already embedded in $\hat{\mathbb{C}}$. This results in flattening of the mesh \mathcal{M} . We denote the flattened set of vertices by $Z = \{z_i\}$.

Generating candidate anti-Möbius transformations. Once the mesh surface is flattened onto the extended plane $\hat{\mathbb{C}}$, it is well known that the entire collection of conformal 1-1 and onto maps $m : \hat{\mathbb{C}} \rightarrow \hat{\mathbb{C}}$ is exactly the Möbius transformations which have the following formula in the extended complex plane:

$$m(z) = \frac{az+b}{cz+d}, \quad (4)$$

where $a, b, c, d \in \mathbb{C}$ satisfying $ad - bc \neq 0$.

We observe that the entire collection of anti-conformal mappings $\bar{m} : \hat{\mathbb{C}} \rightarrow \hat{\mathbb{C}}$ (which contain reflections) can be described by the anti-Möbius transformations:

$$\bar{m}(z) = m(\bar{z}), \quad (5)$$

where m is any Möbius transformation. It is not hard to see from these formulas that both Möbius and anti-Möbius transformations have six (real) degrees of freedom, and fixing the image of three points $z_i \mapsto w_i, i = 1, 2, 3$, on the extended complex plane $z_i, w_i \in \hat{\mathbb{C}}$ uniquely defines either. There are known formulas for prescribing a Möbius transformation from three pairs of points $(z_i, w_i), i = 1, 2, 3$ (see for example [LF09]), however we will suggest a new method for fitting Möbius in the least-squares sense which is applicable if more than three pairs are provided. For prescribing anti-Möbius transformation it is enough to find a Möbius transformation between the pairs $(\bar{z}_i, w_i), i = 1, 2, 3$.

Let $(z_i, w_i), i = 1, 2, \dots, k$ be k pairs of points $z_i, w_i \in \hat{\mathbb{C}}$.

We wish to find a Möbius transformation such that $m(z_i) \approx w_i, i = 1, \dots, k$. Plugging eq. (4), and rearranging the terms we arrive at a system of linear equations:

$$az_i + b - cw_i z_i - dw_i = 0, \quad i = 1, \dots, k \quad (6)$$

An easy approach to solve this system will be to set, say, $a = 1$ and solve the system (6) in the least-squares sense. We might miss a solution with $a = 0$, therefore we can solve again with $b = 1$ and take the solution with the smaller residual error. Note that it is not possible that $a = 0$ and $b = 0$ simultaneously and therefore we cover all Möbius transformations. We note that the suggested least-squares solution minimizes an error based on Euclidean distances with certain non-linear weighting. Therefore setting $a = 1$ or $b = 1$ or $c = 1$ or $d = 1$ will produce different results in general. In practice we are trying all these options and pick the one which minimizes the sum of Euclidean distances squared.

For prescribing generating sets of points, that is, sets of points which uniquely determine an anti-Möbius transformation, we generally need to set at-least three pairs $(z_i, w_i), i = 1, 2, 3$. However, if we restrict ourselves to certain types of symmetry, we can select generating sets in more restrictive sense.

For example, in this paper, we discuss in detail bilateral reflective symmetries. One option is taking a triplet $z_1, z_2, z_3 \in S_1$, and prescribe an anti-Möbius transformation by the requirements that $\bar{m}(z_1) = z_1, \bar{m}(z_2) = z_3, \bar{m}(z_3) = z_2$. If there is a reflective bilateral symmetry $f : \mathcal{M} \rightarrow \mathcal{M}$ then there are such three points z_1, z_2, z_3 that generate that symmetry. Indeed, every bilateral reflective symmetry corresponds to some anti-Möbius transformation \bar{m} . Since, bilateral symmetry satisfies $f^2 = I_d$, then also \bar{m} satisfies the same relation $\bar{m}^2 = I_d$. It means that for every $z_2 \in \hat{\mathbb{C}}$ we can take $z_3 = \bar{m}(z_2)$, and

$$\bar{m}(z_3) = \bar{m}(\bar{m}(z_2)) = z_2.$$

Proposition 7.5 in Section 7 proves that bilateral reflective symmetries have a closed curve of stationary points. Corollary 7.7 then shows that at least two of its points are in the set S_1 . This result justifies taking three points and fixing one to be stationary. Another option for generating candidates for bilateral reflective symmetry is to take four points $z_1, z_2, z_3, z_4 \in S_1$, and find the anti-Möbius that in the least-squares sense (see above) transforms $z_1 \leftrightarrow z_2$ and $z_3 \leftrightarrow z_4$.

In our implementation, we exhaustively check the two types of generating sets (triplets and quadruplets) in S_1 . Furthermore, for numerical stability we first use a anti-Möbius transformation to transform the generating set to a ‘‘canonical frame.’’ In case of three points we simply take an equilateral triangle and map the three points (and the set S_2 with it) to its vertices. In case of four points, we map (as much as possible) to four vertices of a square, trying the two possible different assignments and choosing the one that minimize sum of Euclidean distances squared.

Lastly, we also implemented the Möbius search algorithm

for general symmetry, where four points $z_i \in S_1, i = 1, \dots, 4$, can define a general symmetry by one of the two relations: 1) $z_1 \rightarrow z_2$, and $z_3 \rightarrow z_4$, and 2) $z_1 \rightarrow z_2 \rightarrow z_3 \rightarrow z_4$.

Pruning generating sets. If we consider all possible triplets and quadruplets, there are $\mathcal{O}(|S_1|^4)$ such generating sets. Since S_1 is a small set (usually 5 – 15 points) it is tractable to check all possibilities. However, there are quite a few generating sets which can be detected with high probability as bad, and pruned out. In particular we check to what extent geodesic distance between every two points z, w , which are not mapped to each other via \bar{m} , is preserved: if $\frac{d_g(z, w)}{d_g(\bar{m}(z), \bar{m}(w))} < \delta$ (or the reciprocal is smaller than δ) we prune this set (we use $\delta = 0.9$). Other two useful criteria are: 1) if z, w are mapped to each other, we wouldn't want these two points to be very close to each other: if $d_g(z, w) < \epsilon$ we prune this set, we use $\epsilon = \sqrt{\text{area}(\mathcal{M})/32\pi}$. 2) We check if the AGD gives similar value on points z, w which are mapped to each other: if $\frac{\Phi_{agd}(z)}{\Phi_{agd}(w)} < \gamma$ (or reciprocal rule) we prune this set (we use $\gamma = 0.9$).

Measuring alignment score. The last component of the search procedure is to provide each candidate anti-Möbius transformation with a score, measuring how well it aligns the surface to itself. For that we uniformly sample the surface $S_u \subset \mathcal{M}$ defined by randomizing the first point and iteratively adding the farthest geodesic point from the previously selected set. The anti-Möbius transformation $\bar{m}(z)$ maps the set S_u onto \mathcal{M} , and our goal is to measure how well the two sets $S_u, \bar{m}(S_u)$ align. Note that there is no distance function invariant to anti-Möbius transformations (see [LF09]) and therefore measuring distances (Euclidean or any other distance function) in the flattened space cannot be done in an invariant manner – that is, distances resulted from different alignments are not comparable. Therefore, we define our alignment score based on the *metric* of the surface: we measure geodesic distances between points in $\bar{m}(S_u)$ and S_u . This distance is (by definition) invariant to anti-Möbius transformations. To compute the geodesic distance between pairs of points $z \in \bar{m}(S_u), w \in S_u$ we first need to map z back to \mathcal{M} (we already know where w is). This is done by locating the face of \mathcal{M} containing z and using barycentric coordinates to map back to the mesh. We consider the amount of aligned surface area as our alignment score. We approximate this value by the number of mutually closest (geodesically) pairs of the sets $\bar{m}(S_u), S_u$. Since each point in S_u represents approximately equal area surface patch, the number of mutually closest pairs is an approximation to the amount of area aligned: for example, if the arm of a person is mapped to its belly there will be cluster of points from $\bar{m}(S_u)$ mapped on a sparse set of S_u so there will not be many mutually closest pairs. Note that while our alignment score is generally robust to outliers, it is global since it favors symmetric maps that include the largest possible fraction of a surface.

6. Extracting correspondences

The final step is to take the best anti-Möbius transformation and extract correspondences within the symmetry invariant set $S_2 \subset \mathcal{M}$. This is done in two stages: First, after applying the best anti-Möbius to S_2 and mapping back to the surface we extract geodesically mutually closest pairs $(p_i, q_i), i = 1, 2, \dots, k, p_i, q_i \in S_2$.

Second, we augment them to points in S_2 which did not get assigned. This is done using consistency of geodesic distances as follows. For each point $p \in S_2$ which is not assigned, that is $p \notin \{p_i\}$ we calculate a feature vector

$$\Psi(p) = (d_g(p, p_1), \dots, d_g(p, p_k)). \quad (7)$$

Next, let us define a relative deviation measure

$$D(\Psi(p), \Psi(q)) = \sum_{i=1}^k \left(\frac{d_g(p, p_i) - d_g(q, q_i)}{d_g(p, p_i)} \right)^2.$$

Then, we add to the set of correspondences all pairs (p, q) that have mutually smallest deviation. That is

$$D(\Psi(p), \Psi(q)) = \min_{r \in S_2} D(\Psi(p), \Psi(r)),$$

and symmetrically

$$D(\Psi(q), \Psi(p)) = \min_{r \in S_2} D(\Psi(q), \Psi(r)).$$

Lastly, For every point $p \in \mathcal{M}$ we can use the extracted correspondences $(p_i, q_i), i = 1, 2, \dots, K, p_i, q_i \in S_2$ to find a corresponding point $q \in \mathcal{M}$. This is done by looking for $q \in \mathcal{M}$ which is closest in feature space, that is $D(\Psi(p), \Psi(q)) \leq \min_{r \in S_2} D(\Psi(p), \Psi(r))$.

7. Theory

In this section we provide a rigorous treatment of the relation between intrinsic symmetry and symmetry-invariant functions. We use the theoretical results of this section to construct a stable set of potential intrinsically symmetric points for our algorithm (Section 4). However, we believe these results can be used also in other contexts and actually provide the basis for a direct algorithm which uses the notion of symmetry-invariant functions and sets directly to extract and analyze symmetry properties of surfaces. Throughout this section $f : \mathcal{M} \rightarrow \mathcal{M}$ denotes an intrinsic symmetry.

We start by showing that given a symmetry invariant function Φ , the set $\mathcal{S} \subset \mathcal{M}$ which is defined by the critical points of Φ , $\mathcal{S} = \{p \mid \nabla|_p \Phi = 0\}$, is a symmetry-invariant set that satisfies $f(\mathcal{S}) = \mathcal{S}$. Next, we prove that certain functions (like MGD, AGD and few of its variations) are symmetry-invariant. And lastly, we prove few useful properties regarding stationary points of symmetries and their presence in \mathcal{S} .

First, we establish a connection between gradients of symmetry invariant function at symmetric points:

Proposition 7.1 Let $f : \mathcal{M} \rightarrow \mathcal{M}$ be a symmetry and $\Phi : \mathcal{M} \rightarrow \mathbb{R}$ a symmetry invariant function. Then,

$$\nabla|_{f(p)} \Phi \cdot df|_p = \nabla|_p \Phi, \quad (8)$$

where $df|_p: T_p\mathcal{M} \rightarrow T_{f(p)}\mathcal{M}$ denotes the differential of the map f acting as a linear map between the tangent planes $T_p\mathcal{M}$ and $T_{f(p)}\mathcal{M}$.

Proof Consider the identity $\Phi(f(p)) = \Phi(p)$, and differentiate it. Using the chain rule we get the relation (8). \square

Since f is an isometry its differential is in particular invertible (it is actually orthogonal transformation w.r.t the metric), therefore $\nabla|_p \Phi = 0 \Leftrightarrow \nabla|_{f(p)} \Phi = 0$. This implies that \mathcal{S} satisfies $f(\mathcal{S}) = \mathcal{S}$ (since f^{-1} is also a symmetry).

Next, we prove that the family of AGD functions is symmetry-invariant.

Proposition 7.2 The average geodesic distance function $\Phi_{agd}(\cdot)$ as defined in eq. (1) is invariant to intrinsic symmetries $f: \mathcal{M} \rightarrow \mathcal{M}$. That is,

$$\Phi_{agd}(f(p)) = \Phi_{agd}(p).$$

Furthermore, all the following variations of the average geodesic function are symmetry-invariant:

$$\Phi_{agd}^{\alpha,r}(p) = \int_{d_g(p,q) \leq r} d_g(p,q)^\alpha d\text{vol}_{\mathcal{M}}(q). \quad (9)$$

Proof It is enough to prove the invariance for the general average geodesic distance function $\Phi_{agd}^{\alpha,r}$.

$$\begin{aligned} \Phi_{agd}^{\alpha,r}(f(p)) &= \int_{d_g(f(p),q) \leq r} d_g(f(p),q)^\alpha d\text{vol}_{\mathcal{M}}(q) \\ &= \int_{d_g(p,f^{-1}(q)) \leq r} d_g(p,f^{-1}(q))^\alpha d\text{vol}_{\mathcal{M}}(q) \\ &= \int_{d_g(p,\tilde{q}) \leq r} d_g(p,\tilde{q})^\alpha d\text{vol}_{\mathcal{M}}(\tilde{q}) \\ &= \Phi_{agd}^{\alpha,r}(p), \end{aligned}$$

where in the previous to last equality we have changed the variable $\tilde{q} = f^{-1}(q)$ in the integral and used the fact that f (and f^{-1}) is an isometry and hence volume preserving. \square

Furthermore, the minimal geodesic distance to a symmetry-invariant set \mathcal{S}_1 is also a symmetry-invariant function:

Proposition 7.3 The minimal geodesic distance function $\Phi_{mgd}(\cdot)$, defined in eq. (3), is a symmetry-invariant function.

Proof We directly check this property:

$$\begin{aligned} \Phi_{mgd}(f(p)) &= \min_{q \in \mathcal{S}_1} d_g(f(p),q) \\ &= \min_{q \in \mathcal{S}_1} d_g(p,f^{-1}(q)) \\ &= \min_{q \in \mathcal{S}_1} d_g(p,q) \\ &= \Phi_{mgd}(p), \end{aligned}$$

where the second equality is due to the fact that f^{-1} is also intrinsic symmetry, and the previous to last equality is justified since \mathcal{S}_1 is symmetry-invariant, that is $f^{-1}(\mathcal{S}_1) = \mathcal{S}_1$. \square

So far we have justified the use of the AGD, MGD in defining the symmetry-invariant sets $\mathcal{S}_1, \mathcal{S}_2$ introduced in Section

4. In the rest of this section we would like to prove some properties regarding stationary points of intrinsic symmetries, that is points $p \in \mathcal{M}$ satisfying $f(p) = p$. The key observation is that local analysis of the symmetry at a stationary point provides a representation of the symmetry group as a subgroup of orthogonal matrices acting on \mathbb{R}^2 , i.e., the local tangent plane at the stationary point $T_p\mathcal{M}$.

Theorem 7.4 Let $f: \mathcal{M} \rightarrow \mathcal{M}$ be intrinsic symmetry, and let $p \in \mathcal{M}$ be a stationary point of f , that is $f(p) = p$. Then:

1. The differential $df|_p: T_p\mathcal{M} \rightarrow T_p\mathcal{M}$ is an orthogonal linear transformation, that is,

$$(df|_p)^T (df|_p) = I_d.$$

2. $df|_p$ is orientation preserving iff f is orientation preserving.
3. Let $g: \mathcal{M} \rightarrow \mathcal{M}$ be another symmetry, the map $\Sigma: f \mapsto df|_p$ is a homomorphism. Namely, $\Sigma(f \circ g) = \Sigma(f) \cdot \Sigma(g)$. It means that we can think of the tangent plane $T_p\mathcal{M}$ as the vector space of the representation, and the differential $df|_p$ as the action of the intrinsic symmetry group.
4. The differential $df|_p$ is linear bilateral reflection (w.r.t some direction), iff f is a bilateral reflection. Furthermore, this reflection line (stationary direction) is parallel to the stationary curve of f passing through that point.

Proof First, let us parameterize \mathcal{M} in a neighborhood of p over $T_p\mathcal{M}$. We choose an orthonormal basis in $T_p\mathcal{M}$ and the parametrization can be written as follows

$$\Psi: (x,y) \mapsto (x,y,h(x,y)),$$

where $h(x,y)$ is the height function over the tangent plane and therefore $h(0,0) = 0$, $h_x(0,0) = 0$, and $h_y(0,0) = 0$. The metric tensor in this local coordinate system has the following form:

$$\begin{pmatrix} E & F \\ F & G \end{pmatrix} = \begin{pmatrix} 1 + h_x(x,y)^2 & h_x(x,y)h_y(x,y) \\ h_x(x,y)h_y(x,y) & 1 + h_y(x,y)^2 \end{pmatrix}.$$

In particular at $(x,y) = (0,0)$ we have $E = 1 = G$, $F = 0$. The symmetry f is an isometry fixing p , therefore its differential is a linear map $df|_p: T_p\mathcal{M} \rightarrow T_p\mathcal{M}$ which preserves the metric tensor at $(x,y) = (0,0)$, that is

$$I_d = df^T I_d df,$$

where I_d is the identity matrix. This implies that $df|_p$ is an orthogonal matrix. This proves (1). (2) follows directly from the definition of orientation preserving maps and the fact that isometry is either globally orientation preserving or globally orientation reversing. (3) is an immediate consequence of the chain rule: $d(f \circ g) = df \cdot dg$, using the fact that both $df|_p, dg|_p: T_p\mathcal{M} \rightarrow T_p\mathcal{M}$. Lastly we prove (4). If f is a bilateral reflection with stationary point p , it satisfies $f^2 = I_d$. Using previous argument and taking $g = f$ results in $(df|_p)^2 = I_d$, and $df|_p$ is an orthogonal transformation, and $df|_p$ is orientation reversing, that is $\det(df|_p) < 0$. Therefore the only possible case is that $df|_p$ is a linear bilateral

reflection. On the other direction, if f is not bilateral reflection then $f^2 \neq I_d$ or f is orientation preserving. If $f^2 \neq I_d$ then $(df|_p)^2 \neq I_d$. Otherwise, we can find $q \in \mathcal{M}$ such that $q \neq f^2(q)$. The geodesic from p to q is transformed by f^2 to geodesic from p to $f^2(q)$. However, its tangent vector at $T_p\mathcal{M}$ remain constant since we assume that $df^2|_p = I_d$. This is a contradiction since geodesics are uniquely determined by emanating direction and length. On the other hand, if f is orientation preserving then $\det(df|_p) > 0$ and $df|_p$ cannot be linear bilateral reflection since the eigenvalues of such reflection must be $-1, 1$. Furthermore, the stationary direction of $df|_p$ (eigenvector with eigenvalue 1) has to be parallel to the stationary curve in the case of bilateral reflection. \square

A related fact is that bilateral reflective symmetries necessarily have a closed differentiable curve of stationary points, we prove it for genus zero surfaces:

Proposition 7.5 Bilateral reflective symmetry $f : \mathcal{M} \rightarrow \mathcal{M}$ on genus zero surface has a differentiable closed curve of stationary points.

Proof This proposition follows from the fact that on the extended plane ($\hat{\mathbb{C}}$) f has the form of anti-Möbius \bar{m} , which satisfies the relation $\bar{m} \circ \bar{m} = I_d$, that is anti-involution. This necessarily means (see [Sch79] for more details) that \bar{m} is an inversion w.r.t some circle which it fixes. This circle will correspond to a stationary differentiable closed curve on the surface. \square

Furthermore, every Möbius transformation (on $\hat{\mathbb{C}}$) has at least one fixed point (solution of quadratic equation over the complex numbers). We continue by characterizing the behavior of symmetry invariant functions at stationary points of symmetry.

Theorem 7.6 Let $f : \mathcal{M} \rightarrow \mathcal{M}$ be a symmetry, and $\Phi : \mathcal{M} \rightarrow \mathbb{R}$ a symmetry-invariant function. Then,

1. If f is bilateral reflective symmetry the gradient of Φ is parallel to the curve of stationary points of f .
2. If $f \neq I_d$ is any other symmetry (than bilateral reflection) then any stationary point p is isolated and the gradient of Φ is zero there. In other words, every stationary point is critical point of Φ .

Proof For (1), Proposition 7.5 implies that there is a closed differentiable curve of stationary points of f . Theorem 7.4 tells us that the differential $df|_p$ is a bilateral linear reflection. From Proposition 7.1 we see that at stationary point $p = f(p)$ we have:

$$\nabla|_p \Phi \cdot df|_p = \nabla|_p \Phi. \quad (10)$$

This means that the (transpose of the) differential of f at p fixes the gradient direction of Φ . Therefore, the gradient $\nabla|_p \Phi$ has to be parallel (including the zero vector case) to the reflection line which is the direction of the tangent of the stationary curve. To prove (2) we note that if f has any other symmetry type (and it is not simply identity), Theorem 7.4 implies that $df|_p$ does not have a stationary direction. Fur-

thermore, it is an orthogonal transformation. Eq. (10) now necessarily implies that $\nabla|_p \Phi = 0$. \square

A simple corollary guarantees that the symmetry-invariant sets we construct must contain stationary point(s) if they exist.

Corollary 7.7 Let $f : \mathcal{M} \rightarrow \mathcal{M}$ be an intrinsic symmetry, \mathcal{M} genus zero surface, and $\Phi : \mathcal{M} \rightarrow \mathbb{R}$ symmetry invariant function. Then the set $\mathcal{S} = \{p \mid \nabla|_p \Phi = 0\}$ includes at least one stationary point of f , if it exists. If f is bilateral reflective then \mathcal{S} includes at least two stationary points of f .

Proof If f is not bilateral reflection and has a stationary point, then by Theorem 7.6 this point is critical for Φ and therefore in the set \mathcal{S} . If f is bilateral reflective, it has a differentiable closed stationary curve by Proposition 7.5. Restricted to this curve, Φ achieves maximum and minimum. From Theorem 7.6 the gradient of Φ is parallel to the curve hence it is zero at these two points. These two points are therefore in \mathcal{S} . \square

8. Results

In this section, we report results of experiments with our algorithm for finding intrinsic symmetry correspondences for a wide range of 3D meshes.

Benchmark data sets: The first challenge in evaluating our algorithm is to find a set of meshes for which ground-truth intrinsic symmetry correspondences are known. Unfortunately, to our knowledge, no such data set exists. Images of automatically detected symmetry correspondences have been shown in prior papers (e.g., [BBW*09a, BBW*09b, LTSW09, OSG08, RBB*10]), but we are not aware of any large-scale, quantitative analysis of an intrinsic symmetry algorithm over a wide range of 3D meshes.

To overcome this challenge, we recruited a volunteer to use an interactive program to specify ground truth correspondences $C(T)$ for a large set of points T for all meshes in three data sets (Figure 3):

1. **Watertight:** 400 meshes arranged evenly in 20 object categories, many of which are articulated figures (humans, octopi, four-legged animals, ants, etc.). The meshes were originally created for the SHREC 2007 Watertight Benchmark [GBP07]. Our volunteer selected an average of 10-30 point correspondences per mesh, where some are self-correspondences (stationary points).
2. **SCAPE:** 71 meshes representing a human body in different poses [ASK*05]. All the meshes were fit to scanner data with a common template, and thus they share the same mesh topology. Our volunteer selected 35 point correspondences on one of the meshes, which were then mapped via vertex IDs to all others.
3. **Non-Rigid World:** 75 meshes representing people and animals in a variety of poses [BBK08]. The meshes appear in 8 groups with common topology. Our volunteer selected an average of 15-35 point correspondences from

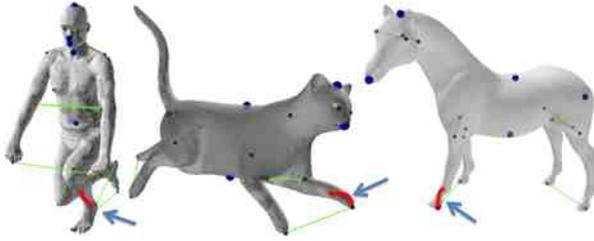


Figure 3: Ground truth examples for SCAPE (human), Watertight’07 (horse), and Non-Rigid World (cat) models. Red curve approximately indicates the length of maximal allowed error for a prediction.

one of each group, which were mapped via vertex IDs to others in the same group.

Note that our volunteer selected semantically meaningful points (rather than automatically extracted ones). For example, points on human body would include tips of hands and legs, top of a head, tip of a nose, etc.

Benchmark evaluation metrics: The second challenge is to develop quantitative metrics to evaluate how well predicted symmetry correspondences matches the ones provided by a human. This is particularly difficult for meshes with partial and high-order symmetries (e.g., k-fold rotations), since multiple different sets of correspondences are all correct, and only one of them has been provided by our volunteer. To avoid this particular problem, we limit our analysis in this paper to meshes where each point has a single global symmetry correspondence (e.g., bilateral reflections) – then the analysis of how well/often symmetric correspondences are found is greatly simplified. This constraint eliminated nine object categories from the Watertight set, leaving humans, glasses, planes, ants, teddy bears, fish, birds, armadillos, busts, pliers, and four-legged animals. All meshes of the other two data sets satisfied this constraint.

For each mesh in the benchmark, we analyze how well the self-mapping f predicted by our algorithm takes every point T_i of the ground-truth set T to its human-labeled correspondence $C(T_i)$ with three metrics:

1. **Geodesic distance:** the average geodesic distance between $C(T_i)$ and $f(T_i)$ for all points in T normalized by the size of the mesh (the square root of the area).
2. **Correspondence rate:** the percentage of points in T for which the geodesic error is less than a distance threshold τ . In our experiments, we set $\tau = \sqrt{\frac{area(\mathcal{M})}{\pi N}}$ (with $N = 20$), such that τ is the radius required for N circular disks to sum up to surface area of the mesh. The size of τ is represented for visual inspection by the length of the red line on each mesh in Figure 3.
3. **Mesh rate:** the percentage of meshes for which the correspondence rate is above a threshold β . In our experiments, we set $\beta = 75\%$.

Benchmark results: Tables 1 and 2 summarize the results

of running the benchmark evaluation on the correspondences predicted by the algorithms proposed in Section 3. Each row of these tables represents a different evaluation metric, and each column represents an average over a different set of meshes.

From these results, we see that the algorithm performs well overall. Looking at the rightmost column of Table 1, we see that at least 75% of the human-specified correspondences were predicted correctly within tolerance for 75% of the meshes. Looking deeper into the results for specific object classes shown in Table 2, we see that our method works very well for objects with little deviation from perfect intrinsic symmetry (e.g., Ants, Birds, etc.), even when they are far from extrinsically symmetric. On the other hand, it performs worst for classes without a small number of clear features (Armadillo), without a single global anti-Möbius transformation mapping the surface onto itself (some busts), and/or where more than one symmetric mapping is almost as good as the correct one (pliers).

	Non-Rigid World	SCAPE Human	SHREC Watertight	All Data Sets
Geodesic	3.3	4.2	1.93	2.65
Corr rate	85%	82%	83%	83%
Mesh rate	76%	72%	75%	75%

Table 1: Evaluation results for three benchmark data sets.

	Human	Glasses	Plane	Ant	Teddy	Plier	Fish	Bird	Armad	Bust	Animal
Geodesic	3.1	1.18	0.94	0.63	2.31	2.84	0.43	0.29	6.47	1.53	1.53
Corr rate (%)	89	88	91	98	86	61	93	99	63	59	83
Mesh rate (%)	90	80	80	100	85	40	95	100	35	50	75

Table 2: Evaluation for classes of SHREC Watertight set.

Representative correspondence predictions produced by our algorithm are shown in Figure 4. In each image, lines are drawn between points predicted as correspondences, and spheres are drawn at positions predicted as stationary points. The main reason for the failure of our algorithm is that some meshes are not isometrically symmetric, and thus there is no single anti-Möbius transformation to align the mesh and its reflection. For each data set, we provide examples sorted by correspondence rates (high to low values). For SCAPE data set we show the worst example marked as success, and the worst example overall, in which the algorithm was able to identify correctly only 22% of the ground truth correspondences (it identified correctly only the bottom part of the mesh).

Comparison to previous work: To evaluate the value of the symmetric point sampling and anti-Möbius search algorithms proposed in this paper, we compare an evaluation of the correspondences produced by our algorithm (left column) to ones achieved with alternative methods in Table 3.

The column titled “Without Symmetric Sampling” represents the results of our method when candidate correspondence points are generated uniformly on the surface (with

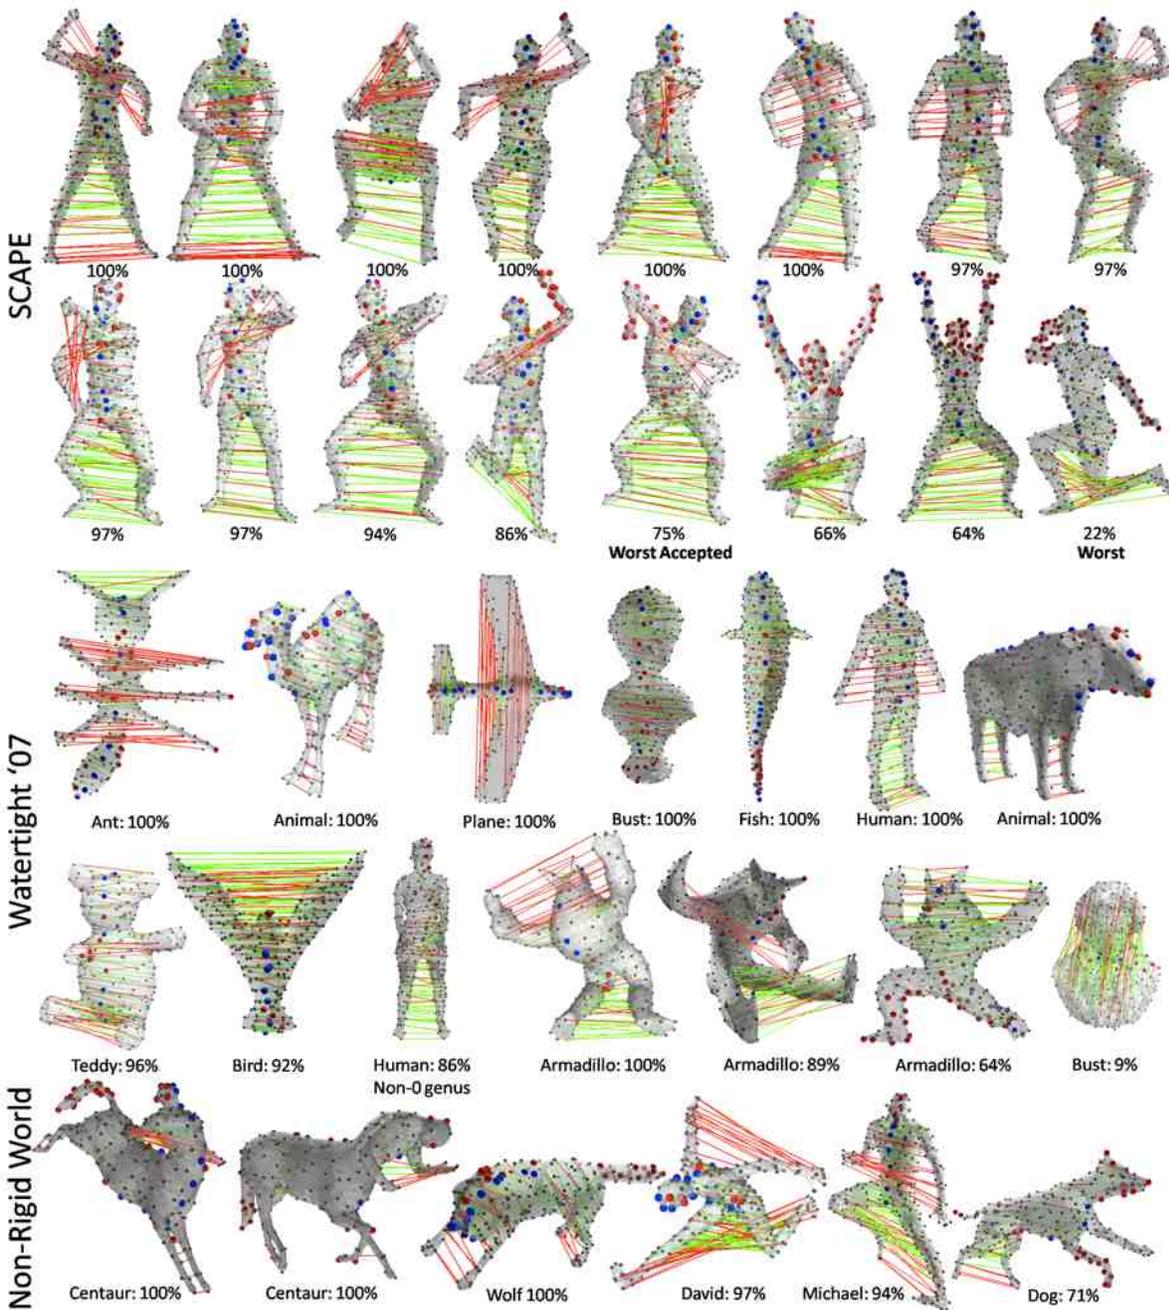


Figure 4: Global Intrinsic Symmetry Benchmark. Percentage below each model indicates the correspondence rate. Green edges and blue points (stationary) come due to mutually closest neighbors in the best anti-Möbius alignment, red edges and vertices come from augmenting the correspondence set.

the furthest point algorithm) instead of the symmetric point sampling algorithm described in Section 4. The symmetric point sampling improves the results slightly, because slight displacement of corresponding predictions to a more symmetric location, does not affect geodesic distances greatly.

The column titled “Mobius Voting (Lipman ‘09)” represents the results that are achieved if the Mobius Voting algorithm were only augmented to work with anti-Möbius transformations without using the improved symmetric point sampling, sampling strategies, alignment score, and correspondence extraction algorithms proposed in this paper. In this case, the difference is quite significant.

	Our Proposed Method	Without Symmetric Sampling	Mobius Voting (Lipman ‘09)
Geodesic	3.49	3.51	6.78
Corr rate (%)	86%	85%	70%
Mesh rate (%)	72%	69%	51%

Table 3: Comparison. Note that our approach performs better than the Möbius Voting approach. Symmetric sampling also slightly improves correspondence results. Contrary to other examples we used 128 samples per mesh due to speed constraints of method Lipman ‘09.

Timing results: For each of the meshes in the benchmark, our algorithm takes 30-90 seconds to run from beginning to end on a 2.2GHz Opteron 275 processor. Specifically, for meshes simplified to 5K triangles, using Dijkstra’s algorithm to compute geodesic distances between all pairs of vertices took approximately 8 seconds; flattening the mesh to the complex plane took 8 seconds; symmetric point sampling took 8 seconds; searching for the best anti-Möbius transformation took 1-60 seconds, depending on the number of points in S_1 (5 points took 1 second, and 17 points took 60 seconds); and, augmenting the correspondence set took 1 second. The main bottlenecks are the memory required for flattening (a sparse matrix is stored in n^2 space in our current implementation, not our contribution) and the time required for searching anti-Möbius transformations (which takes $\mathcal{O}(|S_1|^4)$ in the worst case, but is greatly accelerated by the pruning steps described in Section 3).

Partial and Rotational Symmetries As a final result, we provide examples showing how our algorithm can work for partial and rotational symmetries. Results are not as easily quantifiable in these cases (as described above), and so we resort to depicting our results with images.

Figure 5 (left) shows an example of partial reflectional symmetry detection; the palm of hand is reflectional symmetric, leaving the thumb as an outlier. In Figure 6, we find intrinsic orientation-preserving rotational symmetry.

9. Limitations

In this section we outline the main drawbacks and limitations of our method. First, our method does not handle general

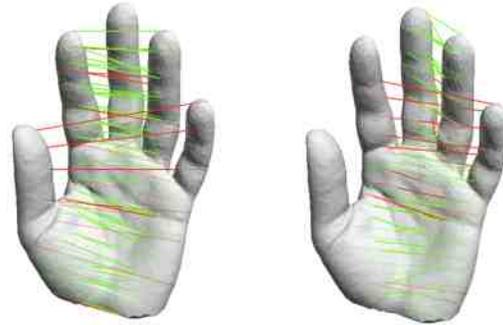


Figure 5: Intrinsically symmetric point correspondences found by our algorithm. Searching for bilateral reflective symmetries, the best Möbius transformation recovered the approximately global reflective symmetry of the hand (left). The second best Möbius recovered partial intrinsic symmetry of the hand, leaving the thumb as an outlier.

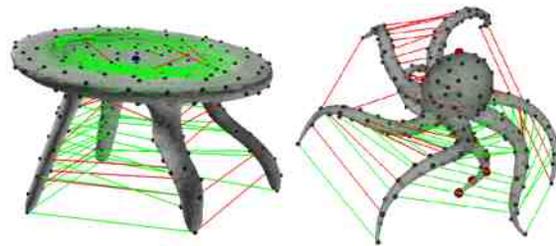


Figure 6: Example of rotational intrinsic symmetry.

partial intrinsic symmetries. While it is relatively robust to outliers (e.g. Figure 5), it searches for conformal mappings that maps most of the surface onto itself. Therefore, it would fail recognizing “small” symmetric parts.

Second, all the theory regarding symmetry invariant functions and sets is provided in the smooth and perfectly symmetric settings. In practice these functions are approximated on meshes that are not perfectly symmetric, and critical points are extracted from these approximations. It is therefore crucial that the symmetry invariant functions possess good approximation properties (i.e., smoothness) – higher regularity of these functions provides more stability. The AGD function for example is defined by integrating geodesic distances and is therefore smooth. The MGD on the other hand, is defined using geodesic distances directly and hence somewhat less regular and stable. Another issue is related to sampling in case of partial symmetry. We note that the partial AGD function (defined and analyzed in Prop. 7.2) would capture partial intrinsic symmetries for $r > 0$ small enough. We have not experimented with this parameter for partial symmetry detection and we consider it as a possible topic for future work.

10. Conclusion and Future Work

This paper investigates using anti-Möbius transformations as a tool for analyzing and detecting intrinsic symmetries in surfaces. It provides two main research contributions: 1) theory and algorithms for construction of symmetric point samples on surfaces, and 2) algorithms to discover intrinsic symmetries by searching the anti-Möbius group efficiently. We find that the combination of these two contributions is effective for finding symmetries in the vast majority of examples in a newly-created, human-labeled benchmark data set for mesh symmetry detection.

This work is a first step towards algorithms that utilize symmetry invariant functions to help identify intrinsic symmetries. We believe that our theoretical findings, such as the symmetry invariant functions, the characterization of intrinsic symmetries as 2D-rigid symmetries at stationary points (group representation), and the results including stationary points in the symmetry invariant sets provide groundwork for further exploring intrinsic symmetry, independently from the approach taken in this work. In particular, one can imagine algorithms for extracting stationary points and/or intrinsic symmetry group description from these tools. One can also think of applications to representing the intrinsic symmetry group as a subgroup of orthogonal transformations in \mathbb{R}^2 . We believe these directions can lead to new algorithms for intrinsic symmetry analysis.

Acknowledgements

This project would not have been possible without the help of people and organizations who provided the data and funding. For the 3D models used in this paper, we thank Daniela Giorgi and AIM@SHAPE for the SHREC 2007 Watertight Models, Drago Arguelov and Stanford University for the SCAPE data set, and Project TOSCA for the Non-rigid World models. For funding in support of this project, we thank the NSERC, NSF (IIS-0612231, CNS-0831374, CCF-0702672, and CCF-0937139), Google, and The Rothschild Foundation.

References

- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: Scape: shape completion and animation of people. *ACM Trans. on Graphics* 24, 3 (2005).
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008.
- [BBW*08] BERNER A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.-P.: A graph-based approach to symmetry detection. In *IEEE/EG Symposium on Volume and Point-Based Graphics* (2008).
- [BBW*09a] BERNER A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.-P.: *Generalized Intrinsic Symmetry Detection*. Tech. Rep. MPI-I-2009-4-005, Max-Planck Institute, August 2009.
- [BBW*09b] BOKELOH M., BERNER A., WAND M., SEIDEL H.-P., SCHILLING A.: Symmetry detection using line features. *Computer Graphics Forum (Eurographics)* 28, 2 (2009).
- [GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: SHREC:SHape REtrieval Contest: Watertight models track, <http://watertight.ge.imati.cnr.it/>, 2007.
- [GCO06] GAL R., COHEN-OR D.: Salient geometric features for partial shape matching and similarity. *ACM Transaction on Graphics* (2006).
- [GY03] GU X., YAU S.-T.: Global conformal surface parameterization. In *SGP '03: Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing* (2003), Eurographics Association, pp. 127–137.
- [LCDF10] LIPMAN Y., CHEN X., DAUBECHIES I., FUNKHOUSER T.: Symmetry factored embedding and distance. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 3 (2010).
- [LE06] LOY G., EKLUNDH J.-O.: Detecting symmetry and symmetric constellations of features. In *ECCV* (2006), pp. 508–521.
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (August 2009).
- [LTSW09] LASOWSKI R., TEVS A., SEIDEL H., WAND M.: A probabilistic framework for partial intrinsic symmetries in geometric data. In *ICCV* (2009).
- [MGP06] MITRA N. J., GUIBAS L., PAULY M.: Partial and approximate symmetry detection for 3d geometry. In *ACM Transactions on Graphics* (2006), vol. 25, pp. 560–568.
- [OSG08] OVSJANIKOV M., SUN J., GUIBAS L.: Global intrinsic symmetries of shapes. *Computer Graphics Forum (Symposium on Geometry Processing)* 27, 5 (2008), 1341–1348.
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2 (1993), 15–36.
- [PSG*06] PODOLAK J., SHILANE P., GOLOVINSKIY A., RUSINKIEWICZ S., FUNKHOUSER T.: A planar-reflective symmetry transform for 3D shapes. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3 (July 2006).
- [RBB*10] RAVIV D., BRONSTEIN A., BRONSTEIN M., KIMMEL R., SAPIRO G.: Diffusion symmetries of non-rigid shapes. In *Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)* (2010).
- [RBBK10] RAVIV D., BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Full and partial symmetries of non-rigid shapes. *International Journal of Computer Vision (IJCV) to appear* (2010).
- [Sch79] SCHWERTDFEGER H.: *Geometry of Complex Numbers*. Courier Dover Publications, 1979.
- [WJG07] WANG S., JIN M., GU X. D.: Conformal geometry and its applications on 3d shape matching, recognition, and stitching. *IEEE Trans. Pattern Anal. Mach. Intell.* 29, 7 (2007), 1209–1220. Member-Yang Wang and Member-Dimitris Samaras.
- [XZT*09] XU K., ZHANG H., TAGLIASACCHI A., LIU L., LI G., MENG M., XIONG Y.: Partial intrinsic reflectional symmetry of 3d shapes. *ACM Transactions on Graphics (SIGGRAPH ASIA)* 28, 5 (2009).
- [ZSCO*08] ZHANG H., SHEFFER A., COHEN-OR D., ZHOU Q., VAN KAICK O., TAGLIASACCHI A.: Deformation-driven shape correspondence. *Comput. Graph. Forum* 27, 5 (2008), 1431–1439.
- [ZWW*10] ZENG Y., WANG C., WANG Y., GU X., SAMRAS D., PARAGIOS N.: Dense non-rigid surface registration using high-order graph matching. In *CVPR '10* (2010).