

Global Intrinsic Symmetries of Shapes

Maks Ovsjanikov Jian Sun Leonidas Guibas

Stanford University

Abstract

Although considerable attention in recent years has been given to the problem of symmetry detection in general shapes, few methods have been developed that aim to detect and quantify the intrinsic symmetry of a shape rather than its extrinsic, or pose-dependent symmetry. In this paper, we present a novel approach for efficiently computing symmetries of a shape which are invariant up to isometry preserving transformations. We show that the intrinsic symmetries of a shape are transformed into the Euclidean symmetries in the signature space defined by the eigenfunctions of the Laplace-Beltrami operator. Based on this observation, we devise an algorithm which detects and computes the isometric mappings from the shape onto itself. We show that our approach is both computationally efficient and robust with respect to small non-isometric deformations, even if they include topological changes.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

1. Introduction

Since most natural objects and phenomena manifest symmetry in some form, detecting and characterizing symmetry is a natural way to understand the structure of shapes. Symmetry provides global information about the structure of the object that is otherwise difficult to capture. Likewise, deviation from perfect symmetry is often an indication of the object's abnormality [RTR*99]. In geometry processing symmetry information has been used for efficient remeshing [PGR07], scan completion [TW05], segmentation [SKS06] as well as shape matching [KFR04].

Despite the importance of symmetry detection in the above tasks, previous work aimed specifically at symmetry detection and classification (e.g. [Rus07a,MGP06]) has concentrated on detecting extrinsic, or embedding-dependent symmetry, with the notable exception of [RBBK07]. Extrinsic symmetry of a shape can be defined as invariance under rigid transformations and possibly scaling. However, because many natural objects are not completely rigid and often undergo non-rigid transformations, this notion of symmetry is very fragile. At the same time, a wide class of deformations, such as articulated motion in humans, preserve the object's internal structure, characterized by the geodesic distances on the object's surface. These deformations leave intact *intrinsic* symmetries of an object, as is defined below.

In this paper we present an algorithm for robustly detecting and categorizing global intrinsic symmetries of shapes. Our main observation is that the Global Point Signatures




Figure 1: Pose-invariant correspondences (intrinsic symmetry) computed with our algorithm

(GPS) defined in [Rus07b] transform intrinsic symmetries on the shape into extrinsic symmetries in a high dimensional signature space. By restricting our attention to a subspace of this signature space we further refine this broad class of extrinsic symmetries to reflectional symmetries around the principal axes. We devise a method for detecting reflectional symmetries in the signature space which allows us to find correspondences on the shape as well as to categorize different symmetries that an object may possess into distinct classes. We also show that our method is robust with respect to small non-isometric deformations and demonstrate the stability of the GPS embedding under topological noise.

2. Related Work

As discussed above, most of the existing work on symmetry detection has concentrated on Euclidean symmetry in two

or three dimensions. Because the search space in this setting is limited to rigid transformations and possibly scaling, methods that optimize over all possible transformations remain efficient. Thus, Podolak and colleagues [PSG*06] develop a transform that measures the reflective symmetries of a shape with respect to all possible planes in a bounding volume, extending the work of Kazhdan and coworkers [KFR04]. Rustamov [Rus07a] augments this transform to include the spatial distribution of the object's asymmetry. Martinet et al. [MSHS06] propose a method for global symmetry detection by analyzing the extrema and spherical harmonic coefficients of generalized moments. Mitra and colleagues [MSHS06] use a voting scheme in the transformation space to detect partial and approximate symmetries of three-dimensional shapes.

In the setting of intrinsic symmetry detection, however, the space of all possible solutions is the set of all endomorphisms on the shape. For this reason it is difficult to devise an efficient transform encoding all possible symmetries or a voting scheme based on point-to-point matches. The problem of intrinsic symmetry detection is thus more closely related to the problem of comparing non-rigid shapes as considered in e.g. [MS05, BBK06]. This latter work proposes a Generalized Multidimensional Scaling (GMDS) approach, which finds a map from points on one shape to points on another minimizing the change in geodesic distances.

In the discrete setting, the problem of intrinsic symmetry detection is analogous to determining if a given graph has a non-trivial automorphism. This problem has unknown complexity between P and NP. Restricted versions, such as fixed-point-free automorphism have been proven to be NP-complete [Lub81].

Also related to our setting is the work by Mitra et al. [MGP07], in which approximate local extrinsic symmetries are used to deform the object to make it more globally extrinsically symmetric, while preserving its structure. However, because the problem is formulated as energy minimization, the resulting deformed object does not necessarily correspond to any intrinsic symmetry present in the original model. For example, a bent object is not intrinsically symmetric, as bending changes the geodesic distances on the surface. To the best of our knowledge, the only work that directly addresses the problem of intrinsic symmetry detection on shapes is [RBBK07], where the authors adapt the GMDS framework to embed the shape onto itself while preserving the geodesic distances as well as possible. Because GMDS is an optimization technique that requires an initial guess, their algorithm proceeds in two stages: first, a coarse initialization is found using local descriptors and geodesic distance histograms. During the second stage, this initialization is used to find a self-mapping on the shape that is as isometric as possible. Our work is different in that we use global descriptors that take the shape's structure into account. As a result, our method is purely algebraic and circumvents the need to

solve a non-convex non-linear optimization problem. Moreover, the descriptors we use also allow us to classify symmetries into discrete classes, which would be difficult to do with the preceding method. Furthermore, as confirmed by our results, correspondences obtained using our algorithm are less sensitive to topological noise.

3. Theoretical Preliminaries

In this section we define the notion of intrinsic symmetry and prove the main results that are required for our method.

3.1. Intrinsic Symmetry

Suppose we are given a compact manifold O without boundary. Following [RBBK07] we call O **intrinsically symmetric** if there exists a homeomorphism $T : O \rightarrow O$ on the manifold that preserves all geodesic distances. That is:

$$g(\mathbf{p}, \mathbf{q}) = g(T(\mathbf{p}), T(\mathbf{q})) \quad \forall \mathbf{p}, \mathbf{q} \in O$$

where $g(\mathbf{p}, \mathbf{q})$ is the geodesic distance between two points on the manifold. In this case, we call the mapping T an intrinsic symmetry.

3.2. Laplace-Beltrami Operator

The main tool that we will use in intrinsic symmetry detection is the Laplace-Beltrami operator Δ (see e.g. [Ros97] for a good introduction), which can be defined entirely in terms of the metric tensor on the manifold independently of the parametrization. The Laplace-Beltrami operator over a compact manifold is bounded and symmetric negative semi-definite. Hence it has an eigen-decomposition, which is a countable set of eigenfunctions $\phi_i : O \rightarrow \mathbb{R}$ and eigenvalues $\lambda_i \in \mathbb{R}$, such that

$$\Delta \phi_i = \lambda_i \phi_i$$

This decomposition has the following useful properties (see Chapter 1 of [Ros97] for the proofs):

1. The eigenfunctions ϕ_i are invariant under isometric deformations of the manifold and are thus intrinsic properties. We address the question of changes under small non-isometric perturbations in Section 4.
2. The set of eigenfunctions forms a complete orthonormal basis for the space of L_2 functions on the manifold. That is, for any square integrable function f :

$$f(\mathbf{p}) = \sum_i a_i \phi_i(\mathbf{p}), \quad \text{where } a_i = \langle f, \phi_i \rangle \quad \forall \mathbf{p} \in O$$

$$\langle \phi_i, \phi_j \rangle = \delta_{ij} = 1 \text{ if } i = j, \text{ and } 0 \text{ otherwise.}$$

3. The Laplacian operator uniquely determines the metric of the manifold and thereby geodesic distances between points. Since the eigendecomposition determines the Laplacian operator, together with property 1. this implies that two manifolds are isometrically deformable into each other if and only if their Laplacian operators have the same eigenvalues and eigenfunctions.

To put the last point more concretely, let O_1 and O_2 be two compact manifolds with Laplace-Beltrami operators Δ_1 and Δ_2 . A homeomorphism $T : O_1 \rightarrow O_2$ preserves all geodesic distances if and only if for every eigenfunction ϕ_i of Δ_2 associated with λ_i , $\phi_i \circ T$ is an eigenfunction of Δ_1 . In particular if T is a self-mapping, we can conclude that T is an intrinsic symmetry if and only if $\phi_i \circ T$ is an eigenfunction associated with eigenvalue λ_i for all ϕ_i .

This immediately leads to an algorithm for detecting intrinsic symmetries:

1. For every bijective mapping T , determine if $\phi_i \circ T$ is an eigenfunction associated with λ_i for all i .
2. If so, T represents an intrinsic symmetry, otherwise such a symmetry does not exist.

This approach unfortunately suffers from the combinatorial complexity of enumerating all bijective mappings on a surface. For this reason, we will look for potential symmetries in the Global Point Signature (GPS) embedding of a manifold introduced by Rustamov in [Rus07b]. Specifically, for each point \mathbf{p} on the manifold O define its signature:

$$s(\mathbf{p}) = \left(\frac{\phi_1(\mathbf{p})}{\sqrt{\lambda_1}}, \frac{\phi_2(\mathbf{p})}{\sqrt{\lambda_2}}, \dots, \frac{\phi_i(\mathbf{p})}{\sqrt{\lambda_i}}, \dots \right). \quad (1)$$

3.3. Main Theorem

Let $s(O)$ denote the embedding of O into the signature space. The main observation of this work is the following Theorem.

Theorem 3.1 *Suppose O has an intrinsic symmetry with an associated isometric self-mapping T . Then $s \circ T \circ s^{-1}$ is an extrinsic symmetry of $s(O)$. In other words, if $\mathbf{p}_1, \mathbf{p}_2 \in O$:*

$$\|s(\mathbf{p}_1) - s(\mathbf{p}_2)\|_2 = \|s(T(\mathbf{p}_1)) - s(T(\mathbf{p}_2))\|_2$$

Proof: By assumption if T is an isometric mapping, and since eigendecomposition is preserved under isometries, $\phi_i \circ T$ must be an eigenfunction associated with eigenvalue λ_i . In other words, $\phi_i \circ T$ must lie in the subspace of the function space spanned by the eigenfunctions ϕ_j associated with λ_i (the number of such ϕ_j is the multiplicity n_i of λ_i). Thus: $\phi_i \circ T = \sum_j a_{ij} \phi_j$. Moreover,

$$\langle \phi_i \circ T, \phi_k \circ T \rangle = \int_O \phi_i(T(\mathbf{p})) \phi_k(T(\mathbf{p})) \text{vol}(\mathbf{p}) = \langle \phi_i, \phi_k \rangle,$$

which equals δ_{ik} by Property 2. The second equality holds because T is required to be a homeomorphism and thus, the integration for both inner products is done over the same set of values, and the volume forms $\text{vol}(\mathbf{p}) = \text{vol}(T(\mathbf{p}))$ since T preserves the metric. Therefore:

$$\sum_j a_{ij}^2 = 1 \text{ and } \sum_j a_{ij} a_{kj} = 0 \text{ if } k \neq i.$$

Because the multiplicity n_i of λ_i is finite [Ros97] p. 32, we can define an orthonormal matrix $A \in \mathbb{R}^{n_i \times n_i}$ such that $A_{ij} = a_{ij}$ as above. Since A is orthonormal, so must be A^T :

$$\sum_i a_{ij}^2 = 1 \text{ and } \sum_j a_{ji} a_{jk} = 0 \text{ if } k \neq i$$




Figure 2: Three cases of eigenfunctions. Blue regions have negative, red have positive, and green have close to zero values. The same eigenfunction ϕ corresponding to a non-repeating eigenvalue, is either (a) positive ($\phi \circ T = \phi$) or (b) negative ($\phi \circ T = -\phi$) depending on the chosen intrinsic symmetry T . (c) An eigenfunction corresponding to a repeating eigenvalue.

Note in case $n_i = 1$, i.e., λ_i is non-repeating, $a_{ii} \in \{+1, -1\}$. Now consider:

$$\begin{aligned} \|s(T(\mathbf{p}_1)) - s(T(\mathbf{p}_2))\|_2^2 &= \sum_{i=1}^{\infty} \left(\frac{\phi_i(T(\mathbf{p}_1)) - \phi_i(T(\mathbf{p}_2))}{\sqrt{\lambda_i}} \right)^2 \\ &= \sum_{k=1}^{\infty} \sum_{i=1}^{n_k} \left(\frac{\phi_{ik}(T(\mathbf{p}_1)) - \phi_{ik}(T(\mathbf{p}_2))}{\sqrt{\lambda_k}} \right)^2, \end{aligned}$$

where the first sum is over all distinct eigenvalues, n_k is the multiplicity of λ_k and ϕ_{ik} is the i^{th} eigenfunction associated with eigenvalue k . Since $\phi_{ik} \circ T = \sum_{j=1}^{n_k} a_{ij} \phi_{jk}$ we have:

$$\begin{aligned} &\sum_{k=1}^{\infty} \sum_{i=1}^{n_k} \left(\frac{\sum_{j=1}^{n_k} a_{ij} (\phi_{jk}(\mathbf{p}_1) - \phi_{jk}(\mathbf{p}_2))}{\sqrt{\lambda_k}} \right)^2 \\ &= \sum_{k=1}^{\infty} \sum_{j,l=1}^{n_k} \sum_{i=1}^{n_k} a_{ij} a_{il} \frac{(\phi_{jk}(\mathbf{p}_1) - \phi_{jk}(\mathbf{p}_2))(\phi_{lk}(\mathbf{p}_1) - \phi_{lk}(\mathbf{p}_2))}{\lambda_k} \\ &= \sum_{k=1}^{\infty} \sum_{j=1}^{n_k} \left(\frac{\phi_{jk}(\mathbf{p}_1) - \phi_{jk}(\mathbf{p}_2)}{\sqrt{\lambda_k}} \right)^2 = \|s(\mathbf{p}_1) - s(\mathbf{p}_2)\|_2^2, \end{aligned}$$

since $\sum_i a_{il} a_{ij} = 1$ if $j = l$ and 0 otherwise. \square

Theorem 2.1 shows that an intrinsic symmetry of an object O induces an extrinsic symmetry (either rotation or reflection) of its GPS embedding. This suggests that we can detect the intrinsic symmetries of an object O by looking for the extrinsic symmetries of $s(O)$ in the signature space.

3.4. Restricted Signature Space

As shown in the proof of Theorem 2.1, the eigenfunctions associated with the repeated eigenvalues can introduce rotational symmetries in the GPS embedding $s(O)$. Since rotational symmetries are notoriously hard to detect in high dimensional spaces, we restrict our attention to the eigenfunctions associated with non-repeated eigenvalues. Note that for the eigenfunctions associated with non-repeating eigenvalues λ_i only one of the following holds:

- $\phi_i \circ T = \phi_i$ we call such ϕ_i *positive*
- $\phi_i \circ T = -\phi_i$ we call such ϕ_i *negative*

In other words, every intrinsic symmetry induces only reflection symmetry around the principal axes in this restricted

signature space. This restriction simplifies our search. Another reason to drop the repeated eigenvalues is that the eigenfunctions associated with them are not stable under small non-isometric perturbations as we show in Section 4. Figure 2 shows that the classification of eigenfunctions associated with non-repeating eigenfunctions is dependent on the intrinsic symmetry T . Ignoring the repeated eigenvalues does not make our method miss any intrinsic symmetries, but may mix different isometric self-mappings, i.e., some symmetries become un-distinguishable by only considering the non-repeated eigenvalues. In general, a further refinement step may be needed as discussed in Section 3.

Note that in general the presence of repeated eigenvalues is neither necessary, nor sufficient for intrinsic symmetries. It is possible to construct graphs with no automorphisms, whose Laplacians have repeated eigenvalues, and conversely automorphic graphs whose Laplacians do not have repeated eigenvalues. We conjecture that the same holds for smooth manifolds. It is possible, however, to prove that the Laplace-Beltrami operator on a manifold with an intrinsic symmetry T such that $T \circ T \neq Id$, or with two intrinsic symmetries S, T such that $S \circ T \neq T \circ S$, must have repeated eigenvalues. We omit these proofs for the lack of space.

We also consider the projection of the signature space onto its first d components corresponding to the first d non-repeating eigenvalues. Increasing d results in higher discriminative power (points will have more unique signatures) at the cost of increased complexity and exposure to numerical errors. Moreover, the eigenvectors associated with the smallest eigenvalues are more stable under non-isometric deformations, and approximately preserve the positive/negative properties discussed above. Therefore, using fewer eigenfunctions results in a more robust method at the cost of accuracy of differentiating between points.

In summary, for each point \mathbf{p} on the manifold O , we define its *restricted* GPS signature $s(\mathbf{p})$ to be:

$$s(\mathbf{p}) = \left(\frac{\phi_1(\mathbf{p})}{\sqrt{\lambda_1}}, \frac{\phi_2(\mathbf{p})}{\sqrt{\lambda_2}}, \dots, \frac{\phi_d(\mathbf{p})}{\sqrt{\lambda_d}} \right), \quad (2)$$

where λ_i is non-repeated for any $1 \leq i \leq d$.

4. Symmetry Detection and Classification

As discussed above, given an intrinsic symmetry T , in the restricted signature space each eigenfunction is either positive or negative. Thus every intrinsic symmetry T can be identified with a sequence of signs (+ or -). That is, given a manifold O with an intrinsic symmetry T , in the restricted signature space T induces a sequence of signs determined by whether $\phi_i \circ T = \phi_i$ or $-\phi_i$.

Our approach to detecting an intrinsic symmetry T will be to recover the sequence of signs induced by T and use it to find point-to-point correspondences. Note that the restriction of the signature space to non-repeating eigenvalues will




Figure 3: T is a cyclic intrinsic symmetry with order 5, and thus none of the eigenfunctions can be negative with respect to it. A negative eigenfunction would imply $f(\mathbf{p}) = -f(T^5(\mathbf{p})) = -f(\mathbf{p}) \forall \mathbf{p}$

not make us miss any intrinsic symmetries. This is because every intrinsic symmetry is associated with a sequence of signs in the restricted signature space regardless of d . This can also be seen in the proof of Theorem 2.1, as the result holds even if $s(\mathbf{p})$ is defined over a finite number of distinct non-repeating eigenvalues. Nevertheless, it is possible that different intrinsic symmetries induce the same sequence of signs. For example in Figure 3, since the symmetry T does not admit any negative eigenfunctions, it induces a sequence of all pluses. So do T^2, T^3, T^4 and $T^5 = Identity$. Imagine in the extreme case of torus, there is an infinite number of symmetries along the longitude, which all induce the same sequence of signs.

In general, recovering the correct sequence of signs induced by T reduces the search space for T but may not actually determine T . For example, in the case of the Star model in Figure 3, the sequence of all pluses reduces the set of potential correspondences down to 4. That is, given a query point \mathbf{q} , there are 4 other points having the same signature in the restricted signature space. In this work, we assume that the recovered sequence of signs has enough precision to detect its inducing intrinsic symmetry. For most of the models used in our experiments, this assumption was fulfilled. We present a simple way to relax this assumption in Section 3.3 and leave the problem of disambiguating all intrinsic symmetries as an interesting problem for future work. Nevertheless, even when two intrinsic symmetries induce the same sign sequence, our method can be used to reduce the set of candidate correspondences for each point down to two.

Given a closed manifold O , suppose it has an intrinsic symmetry $T : O \rightarrow O$. Let $s_i(\mathbf{q})$ be the i^{th} component of the restricted GPS signature of $\mathbf{q} \in O$ as in (2) and $S = (S_1, \dots, S_d)$ be the sign sequence corresponding to T . That is, $S_i = 1$ if $T \circ \phi_i = \phi_i$ and -1 if $T \circ \phi_i = -\phi_i$. Since the signatures are only defined on the non-repeating eigenvalues, we must have, for any point $\mathbf{q} \in O$:

$$s(T(\mathbf{q})) = (S_1 s_1(\mathbf{q}), S_2 s_2(\mathbf{q}), \dots, S_d s_d(\mathbf{q})) \quad (3)$$

Note, in particular, $|s(T(\mathbf{q}))_i| = |s(\mathbf{q})_i| \forall i$. Conversely, given a sign sequence S , we can test whether this sequence is associated with any intrinsic symmetry on the manifold by modifying the signature s of each point \mathbf{q} to $s'(\mathbf{q}) =$

$(S_1 s_1(\mathbf{q}), S_2 s_2(\mathbf{q}), \dots, S_d s_d(\mathbf{q}))$ and considering:

$$E(S) = \sum_{\mathbf{q} \in O} \min_{\mathbf{p} \in O} \|s'(\mathbf{q}) - s(\mathbf{p})\|_2^2$$

For a sign sequence S corresponding to an intrinsic symmetry, $E(S) = 0$ because for every point there must exist a corresponding symmetric point. In the case of approximate symmetry, $E(S)$ must be small because non-repeating eigenfunctions are continuous and, at least in the discrete setting, change continuously under non-isometric deformations (Section 4). Note that $E(S) = 0$ is necessary but may not be sufficient for the sign sequence to correspond to an intrinsic symmetry. In practice, we observe that this does not happen for moderately large d (15 in our experiments),

Since the signature $s(\mathbf{q})$ is defined on d eigenfunctions corresponding to non-repeating eigenvalues, a simple algorithm to retrieve the candidate sign sequences would be to test all 2^d possibilities and rank them by the error $E(S)$. This method works well. For instance, for the Female model in Figure 4, the three combinations of lowest error correspond to the bilateral symmetry, the front/back symmetry and their composition respectively. However, the complexity of this method is exponential in d , the number of components chosen for the GPS embedding. In order to overcome this limitation, we propose the following procedure to detect potentially negative eigenfunctions.

4.1. Detecting Negative Eigenfunctions

As is shown on Figure 2 the classification of eigenfunctions into negative and positive ones depends on a particular intrinsic symmetry T . On the other hand, there can be eigenfunctions that are positive with respect to all intrinsic symmetries, and detecting those will allow us reduce the search space for the correct sign sequences. In practice, roughly half of the eigenfunctions are always positive, so it is beneficial to identify them. We propose the following simple procedure.

1. Given an eigenfunction ϕ_i , for each point \mathbf{p} , let

$$\begin{aligned} s_i^+(\mathbf{p}) &= (|s_1(\mathbf{p})|, |s_2(\mathbf{p})|, \dots, s_i(\mathbf{p}), \dots, |s_d(\mathbf{p})|) \text{ and} \\ s_i^-(\mathbf{p}) &= (|s_1(\mathbf{p})|, |s_2(\mathbf{p})|, \dots, -s_i(\mathbf{p}), \dots, |s_d(\mathbf{p})|). \end{aligned}$$

That is, each component of $s_i^+(\mathbf{p})$ and $s_i^-(\mathbf{p})$ is the same, and equals the absolute value of the corresponding component of $s(\mathbf{p})$, except their i^{th} component, which is equal and negative of the i^{th} component of $s(\mathbf{p})$ respectively.

2. Assume there exists an intrinsic symmetry T for which ϕ_i is a negative eigenfunction. Then for every point \mathbf{p} there must exist a corresponding point \mathbf{q} such that $s_i^+(\mathbf{q}) = s_i^-(\mathbf{p})$. Thus, we compute:

$$E_i = \sum_{\mathbf{p} \in O} \min_{\mathbf{q} \in O} \|s_i^-(\mathbf{p}) - s_i^+(\mathbf{q})\|_2^2,$$

which should be 0. In practice, we consider the i^{th} eigenfunctions potentially negative if E_i is small.

4.2. Coupling

Given a list of potentially negative eigenfunctions we would like to partition them into groups such that all of the functions within a group are negated by the same symmetry. For this, we follow a similar procedure as above. Specifically to check if the i^{th} eigenfunction ϕ_i and the j^{th} eigenfunction ϕ_j (both potentially negative) are coupled, we consider the following signature for each point $\mathbf{p} \in O$

$$\begin{aligned} s_{ij}^+(\mathbf{p}) &= (|s_1(\mathbf{p})|, \dots, s_i(\mathbf{p}), \dots, s_j(\mathbf{p}), \dots, |s_d(\mathbf{p})|) \text{ and} \\ s_{ij}^-(\mathbf{p}) &= (|s_1(\mathbf{p})|, \dots, s_i(\mathbf{p}), \dots, -s_j(\mathbf{p}), \dots, |s_d(\mathbf{p})|). \end{aligned} \quad (4)$$

That is, each component of $s_{ij}^+(\mathbf{p})$ and $s_{ij}^-(\mathbf{p})$ is the same, and equals the absolute value of the corresponding component of $s(\mathbf{p})$, except for their i^{th} and j^{th} components. Then we compute:

$$E_{ij} = \sum_{\mathbf{p} \in O} \min_{\mathbf{q} \in O} \|s_{ij}^-(\mathbf{p}) - s_{ij}^+(\mathbf{q})\|_2^2.$$

If ϕ_i and ϕ_j are negative eigenfunction for different symmetries then $\phi_i \circ T = \phi_i$ and $\phi_j \circ T = -\phi_j$, resulting in $E_{ij} = 0$ or $E_{ji} = 0$. Otherwise, E_{ij} and E_{ji} must both be large, and in particular, larger than E_i or E_j . In this case, we call ϕ_i and ϕ_j are coupled.

Once the pairwise coupling information is computed, we group the potentially negative eigenfunctions as follows. Starting with a list of pairs of coupled negative eigenfunctions, we join two pairs if they have a non-empty intersection (if at least one eigenfunction is in both pairs). We then iterate this procedure until we have disjoint sets of negative eigenfunctions. Finally, we treat each such set as a sign sequence S , by setting $S_i = -1$ if ϕ_i is in the set, and 1 if it is not. This results in the list of sign sequences, that we use to compute point-to-point correspondences within the shape.

4.3. Point-To-Point Correspondences

Using the sign sequences obtained by the procedure above, we look for correspondences in a straightforward fashion. Given a query point \mathbf{q} , we go through each sequence of signs. Note under our assumptions, distinct sequences of signs correspond to a distinct intrinsic symmetry. For a sequence of sign $S = (S_1, \dots, S_d)$, define $s'(\mathbf{p}) = (S_1 s_1(\mathbf{q}), S_2 s_2(\mathbf{q}), \dots, S_d s_d(\mathbf{q}))$ as in Eqn (3). We then look for the closest point

$$\min_{\mathbf{p} \in O} \|s'(\mathbf{q}) - s(\mathbf{p})\|,$$

and declare \mathbf{p} , the corresponding point to \mathbf{q} . Since an object can have multiple intrinsic symmetries, by going through each sign sequence, we can determine different corresponding points.

It is also possible, albeit rare, that two intrinsic symmetries correspond to the same sign sequence in the restricted signature space. Then, for a given query point \mathbf{q} and sign sequence, there are two corresponding points \mathbf{p}_1 and \mathbf{p}_2 . A

simple way to solve this problem would be to look for a non-repeating eigenfunction ϕ_i such that $\phi_i(\mathbf{p}_1) \neq \phi_i(\mathbf{p}_2)$. Including ϕ_i in the definition of the restricted signature, will ensure that the two intrinsic symmetries $T(\mathbf{q}) = \mathbf{p}_1$ and $S(\mathbf{q}) = \mathbf{p}_2$ will correspond to different sign sequences. If all eigenvalues are distinct, such ϕ is guaranteed to exist.

5. Implementation

To implement the method described above, we use triangular meshes to approximate smooth manifolds. Several discrete schemes have been proposed to approximate the Laplace-Beltrami operator on triangular meshes. One widely used discrete Laplace operator is the cotangent (COT) scheme, originally proposed by Pinkall and Polthier [PP93]. Recently, Belkin et al [BSW08] proposed a scheme based on the heat equation, which is proven to converge point-wise on arbitrary meshes. Although it is well known that no discrete Laplace operator can share all of the properties of its continuous counterpart [WMKG07], in our experiments these schemes produce eigenfunctions that approximately preserve the positive/negative properties discussed above. We base our computation on [BSW08] since we observe it gives more stable eigenfunction computation. As stated in [Rus07b], the Laplacian matrix L can be rewritten as $S^{-1}M$ where S is a diagonal and M is symmetric. We present a stability argument on solving the eigen-problem $L\phi = \lambda\phi$ before we describe the implementation and its complexity.

Stability Argument. For our approach to practically work, we need to show the eigenvectors are stable under small non-isometric perturbations. Denote \tilde{O} to be a small non-isometric perturbation of object O . Let \tilde{D} be a mesh that discretizes object \tilde{O} . We move each vertex of \tilde{D} onto its closest point on O and obtain a new mesh that discretizes object O . We denote this new mesh as D . Since \tilde{O} is just a small perturbation of O , we assume D well approximates O . Furthermore, if we let $L = S^{-1}M$ and $\tilde{L} = \tilde{S}^{-1}\tilde{M}$ denote the Laplacian matrices of D and \tilde{D} respectively, we have that $|\delta M = M - \tilde{M}|$ and $|\delta S = S - \tilde{S}|$ are both small. We assume that both D and \tilde{D} are dense enough and well approximate O and \tilde{O} respectively. Let λ_i and ϕ_i denote the i^{th} eigenvalue and eigenvector of matrix L , and $\tilde{\lambda}_i$ and $\tilde{\phi}_i$ are the i^{th} eigenvalue and eigenvector of matrix \tilde{L} .

Based on matrix perturbation theory, up to the first order of δM and δS , we have

$$\begin{aligned}\tilde{\lambda}_i &= \lambda_i + \phi_i^T([\delta M] - \lambda_i[\delta S])\phi_i \quad \text{and} \\ \tilde{\phi}_i &= \phi_i(1 - \frac{1}{2}\phi_i^T[\delta S]\phi_i) + \sum_{j \neq i} \frac{\phi_j^T([\delta M] - \lambda_i[\delta S])\phi_i}{\lambda_i - \lambda_j} \phi_j.\end{aligned}$$

From these two expressions, we conclude that the eigenvectors with spread out eigenvalues are more stable against small perturbations. Note that in general, eigenvectors corresponding to repeating eigenvalues are unstable.

Since we choose stable eigenvectors, our algorithm can be

applied to **approximately symmetric** objects. An object is approximately symmetric if it is close to some intrinsically symmetric object. Here we consider two objects close if the points of one can be perturbed onto the points of the other, as discussed above.

Complexity. Our algorithm consists of two main steps: computing the eigenvectors of Laplacian matrix to build the restricted GPS embedding, and finding nearest neighbors in the restricted signature space. In the experiments below we use the sparse eigen solver implemented in the **ARPACK** package to compute the restricted GPS embedding. This allows us to compute the eigenvectors of meshes with up to 40,000 points. Note that [VL08] presents an out-of-core algorithm that finds an eigendecomposition of the Laplace-Beltrami operator on meshes with up to a million points.

For nearest neighbor search, we use the kd-tree data structure implemented in the **ANN** library. Note that the complexity of a nearest neighbor query in the kd-tree has been observed to depend on the underlying dimensionality of the data more than on the dimension d of the embedding space [Moo91], and equals $O(\log n)$ for a data-set of fixed underlying dimension. Since the GPS embedding is a homeomorphism, the underlying dimensionality is fixed (equals 2 for surfaces in 3D), regardless of d . The complexity of determining whether an eigenfunction is negative and the overall complexity of our method are approximately $O(dn \log n)$ and $O(d^3 n \log n)$, respectively. In our experiments, we use the first 15 non-repeated eigenvectors, i.e., $d = 15$ unless mentioned otherwise. Repeated eigenvalues were not present in any model, except for the Female model, where one pair was removed using conservative scheme based on the relative difference between eigenvalues. Table 1 shows the timing of our algorithm on the models shown in the result section.

Model	#vertices	#components	GPS	Cor
FEMALE	8968	15	22	67
		30	24	361
MAN	14603	15	74	143
		30	78	786
HIP	19979	15	131	366
		30	138	2980
KNOTS	38400	15	464	573
		30	498	4346

Table 1: Timing in seconds for the different stages on a PC with Quad-Core Intel Xeon 5300 and 3GBytes main memory.

6. Results

Figure 4 shows the point-to-point correspondences obtained with our algorithm. In colored models, we show a positive eigenfunction, in which the same colors roughly represent corresponding regions under the intrinsic symmetry. Note that although we only show correspondences for a few points, our algorithm generates a correspondence for every point. Our method captures Euclidean symmetry when it is present such as the reflection of the Hip model and the rotation of the Twirl model. The Moebius strip with the follow-

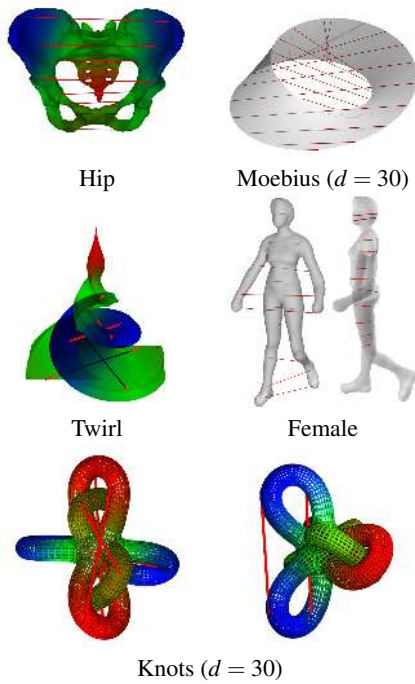


Figure 4: Correspondences obtained with our algorithm on different models.

ing parametrization:

$$p(u, v) = \begin{pmatrix} \cos(u) + v \cos(u/2) \cos(u) \\ \sin(u) + v \cos(u/2) \sin(u) \\ v \sin(u/2) \end{pmatrix}$$

has an intrinsic symmetry, which can be found analytically: $T((x, y, z)) = (x, -y, -z)$. On a discrete model, our algorithm generates exactly this map. The knots model also possesses one intrinsic symmetry, which is not easy to observe. Our algorithm computes the symmetry map as shown in Figure 4 from two different viewing angles. For the Female model, our algorithm detects two separate approximate symmetries: bilateral and front/back.

Deformation. We also tested our algorithm on two sets of deformed models. The three human poses shown on the first row of Figure 5 were taken from the SCAPE dataset, which is based on scans of a real person. For these poses, our method correctly captures the intrinsic bilateral symmetry of the human. The poses shown on the second row of Figure 5, were obtained by the method in [AOW*08] which produces as-rigid-as possible volume preserving deformations. Starting with an extrinsically symmetric model, we deformed it, and computed the maximum change in the geodesic distances between any pair of points on the model using the fast-marching method. The numbers under the models show this geodesic distortion of the deformation. As can be seen, although these deformations are not isometric, we are able to detect approximate intrinsic symmetries and compute consistent correspondences as long as the distortion is not very

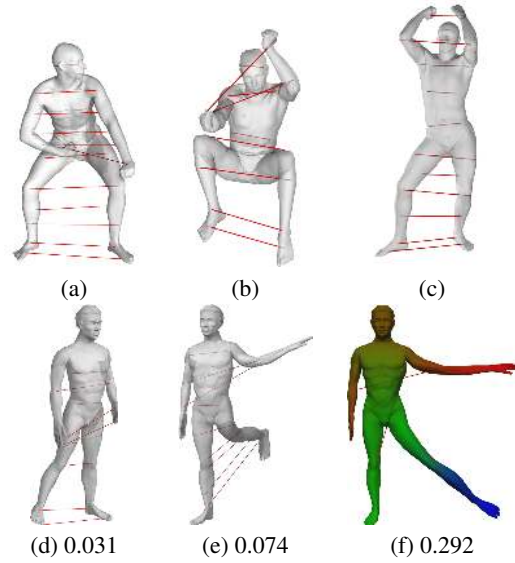


Figure 5: (a-c) Correspondences on models taken from the SCAPE dataset. (d-f) Correspondences on a model undergoing as-rigid-as possible deformations. The numbers represent the geodesic distortion of the deformation. Our method produces consistent correspondences until a large distortion (f). In this case the first non-repeated eigenfunction (shown) fails to be either positive or negative.

large. When the model undergoes a large non isometric deformation Figure 5 (f), our method fails as the first non-repeated eigenfunction is neither positive nor negative.

Topological Change. Rustamov [Rus07b] pointed out that GPS embedding is stable against topological noise but did not give an explanation. Certainly, the GPS embedding of an object changes as its topology changes as the eigenfunctions and the eigenvalues characterize objects up to isometry, which is a topological invariant. However we observe that the introduction of small topological “short-cuts” only induce changes in the high frequency eigenfunctions corresponding to large eigenvalues as shown in Figure 6. This observation explains why the restricted GPS embed-




Figure 6: Boy, (a) a small handle added between two fingers. (b) a larger handle added between the right arm and the torso. The plots show the changes of first two 200 eigenfunctions (excluding the repeated ones).

ding is stable against topological noise since the eigenfunctions of high frequency are scaled down proportional to their eigenvalues and truncated early. This enables us to detect intrinsic symmetries even with small topological noise, when




Figure 7: Correspondences before and after the topological changes shown in Figure 6.

geodesic-based methods experience problems [RBBK07]. Figure 7 shows the result of our algorithm applied on the models with small topological changes.

7. Conclusion and Future Work

We have described an algorithm to detect intrinsic symmetries and extract point correspondences on the shape. Our main observation is that Global Point Signatures defined in [Rus07b] map intrinsic symmetries into extrinsic symmetries in the signature space. We then observe that considering only non-repeating eigenvalues reduces this space of extrinsic symmetries in high dimension to reflectional symmetries, which can be detected efficiently. We demonstrated the applicability of our method by computing correspondences on a set of objects with approximate symmetries and showed that they are robust with respect to small non-isometric deformations and topological noise.

As noted in Section 3, it is possible that multiple intrinsic symmetries induce the same sign sequence. In the future, we are planning to add a refinement step to our framework, that would allow to distinguish different symmetries if they induce the same sign sequence. In addition to the method in Section 3.3, one possibility is to use spectral clustering techniques described in [LH05] to find the clusters of geodesically consistent correspondences. We are also considering ways to extend our method to incorporate the repeated eigenfunctions to achieve more robustness. Furthermore, although our method can be used to classify different classes of intrinsic symmetries, we would like to find a compact representation of all intrinsic symmetries, including continuous symmetries. Finally, we are working to extend our method to manifolds with boundaries and local and partial intrinsic symmetries.

Acknowledgments This work was supported by NSF grants ITR 0205671 and FRG 0354543, NIH grant GM-072970, DARPA grant HR0011-05-1-0007, a grant from Agilent Corporation, and the Max-Planck Center for Visual Computing and Communication. The authors would also like to thank Alex and Michael Bronstein and the anonymous reviewers for the helpful suggestions and comments.

References

[AOW*08] ADAMS B., OVSJANIKOV M., WAND M., SEIDEL H.-P., GUIBAS L.: Meshless modeling of deformable shapes and their motion. *To appear in ACM SIGGRAPH Symposium on Computer Animation* (2008).

- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. National Academy of Sciences (PNAS)* 103, 5 (January 2006), 1168–1172.
- [BSW08] BELKIN M., SUN J., WANG Y.: Discrete laplace operator on meshed surfaces. In *SoCG* (2008), p. to appear.
- [KFR04] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Symmetry descriptors and 3d shape matching. In *Proceedings of SGP* (2004), pp. 115–123.
- [LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *Proceedings of ICCV* (2005), pp. 1482–1489.
- [Lub81] LUBIW A.: Some np-complete problems similar to graph isomorphism. *SIAM Journal on Computing* 10, 1 (1981), 11–21.
- [MGP06] MITRA N. J., GUIBAS L., PAULY M.: Partial and approximate symmetry detection for 3d geometry. In *ACM Trans. Graph.* (2006), vol. 25(3), pp. 560–568.
- [MGP07] MITRA N. J., GUIBAS L., PAULY M.: Symmetrization. In *SIGGRAPH* (2007), vol. 26(3), pp. #63, 1–8.
- [Moo91] MOORE A.: *An introductory tutorial on kd-trees*. Tech. Rep. 209, University of Cambridge, 1991.
- [MS05] MÉMOLI F., SAPIRO G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.* 5, 3 (2005).
- [MSHS06] MARTINET A., SOLER C., HOLZSCHUCH N., SILLION F. X.: Accurate detection of symmetries in 3d shapes. *ACM Trans. Graph.* 25, 2 (2006), 439–464.
- [PGR07] PODOLAK J., GOLOVINSKIY A., RUSINKIEWICZ S.: Symmetry-enhanced remeshing of surfaces. In *Symposium on Geometry Processing* (July 2007).
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 1 (1993), 15–36.
- [PSG*06] PODOLAK J., SHILANE P., GOLOVINSKIY A., RUSINKIEWICZ S., FUNKHOUSER T.: A planar-reflective symmetry transform for 3D shapes. *SIGGRAPH* 25, 3 (2006).
- [RBBK07] RAVIV D., BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Symmetries of non-rigid shapes. In *NRTL* (2007), pp. 1–7.
- [Ros97] ROSENBERG S.: *Laplacian on a Riemannian manifold*. Cambridge University Press, 1997.
- [RTR*99] REITE M., TEALEA P., ROJASA D. C., SHEEDERA J., ARCINIEGASA D.: Schizoaffective disorder: evidence for reversed cerebral asymmetry. *Biological Psychiatry* 46, 1 (1999).
- [Rus07a] RUSTAMOV R. M.: Augmented symmetry transforms. In *Proceedings of SMI* (2007), pp. 13–20.
- [Rus07b] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of SGP* (2007), pp. 225–233.
- [SKS06] SIMARI P., KALOGERAKIS E., SINGH K.: Folding meshes: hierarchical mesh segmentation based on planar symmetry. In *Proceedings of SGP* (2006), pp. 111–119.
- [TW05] THRUN S., WEGBREIT B.: Shape from symmetry. In *Proceedings of ICCV* (2005), pp. 1824–1831.
- [VL08] VALLET B., LEVY B.: Manifold harmonics. *Computer Graphics Forum* 27(2) (*Proc. Eurographics*) (2008).
- [WMKG07] WARDETZKY M., MATHUR S., KÄLBERER F., GRINSPUN E.: Discrete laplace operators: no free lunch. In *Proceedings of SGP* (2007), pp. 33–37.