

Multi-Scale Partial Intrinsic Symmetry Detection

Kai Xu^{*†} Hao Zhang[‡] Wei Jiang^{*} Ramsay Dyer[§] Zhiquan Cheng^{*} Ligang Liu^{*} Baoquan Chen[†]
^{*}National University of Defense Technology [†]Shenzhen VisuCA Key Lab/SIAT
[‡]Simon Fraser University [§]INRIA, GEOMETRICA ^{*}University of Science and Technology of China




Figure 1: Multi-scale partial intrinsic symmetry detection: five symmetry scales (large to small) are detected. Each symmetric region is shown in uniform color. Note the detection of inter- and intra-object symmetries, as well as cylindrical symmetry of the limbs.

Abstract

We present an algorithm for *multi-scale* partial intrinsic symmetry detection over 2D and 3D shapes, where the scale of a symmetric region is defined by intrinsic distances between symmetric points over the region. To identify prominent symmetric regions which overlap and vary in form and scale, we decouple scale extraction and symmetry extraction by performing two levels of clustering. First, significant symmetry scales are identified by clustering sample point *pairs* from an input shape. Since different point pairs can share a common point, shape regions covered by points in different scale clusters can overlap. We introduce the *symmetry scale matrix* (SSM), where each entry estimates the likelihood two point pairs belong to symmetries at the same scale. The *pair-to-pair* symmetry affinity is computed based on a pair signature which encodes scales. We perform spectral clustering using the SSM to obtain the scale clusters. Then for all points belonging to the same scale cluster, we perform the second-level spectral clustering, based on a novel *point-to-point* symmetry affinity measure, to extract partial symmetries at that scale. We demonstrate our algorithm on complex shapes possessing rich symmetries at multiple scales.

Links: [DL](#) [PDF](#) [WEB](#) [DATA](#)

1 Introduction

Symmetry is ubiquitous in nature and in manufactured artifacts. The study of shape symmetry has attracted much attention in com-

puter graphics lately [Mitra et al. 2012]. Most existing works focus on detecting extrinsic symmetries and a popular approach is transformation space voting [Mitra et al. 2006]. In the intrinsic case, most efforts are devoted to global symmetries [Ovsjanikov et al. 2008; Lipman et al. 2010]. It is generally recognized that partial intrinsic symmetry detection is a challenging problem since it has to deal with larger search spaces for both the symmetric regions (compared to global symmetry analysis) and the symmetry-revealing transforms (compared to extrinsic symmetry detection).

A partial intrinsic symmetry over a shape is a subregion with associated self-homeomorphisms that preserve all pairwise intrinsic distances [Mitra et al. 2012]. In this paper, we address the problem of partial intrinsic symmetry detection over 2D and 3D shapes. More importantly, our goal is to detect such symmetries at *multiple scales*, where we define the scale of a symmetric region based on the intrinsic distances between symmetric points over the region.

Complex shapes often exhibit multiple symmetries that overlap and vary in form and scale (see Figure 1). Multi-scale analysis enables the construction of high-level, coarse-to-fine representations for such shapes [Wang et al. 2011] to improve shape understanding and facilitate solutions to such problems as shape correspondence, editing, and synthesis. However, multi-scale symmetry analysis poses additional challenges. The problem is at first complicated by a new search dimension, the scale dimension. Also, existing approaches to intrinsic symmetry detection including those based on region growing [Xu et al. 2009], partial matching [Raviv et al. 2010], and symmetry correspondence [Lipman et al. 2010], cannot extract symmetries that physically overlap.

Symmetry detection inherently involves a grouping of shape elements deemed to be symmetric to each other. A clustering approach often facilitates the detection of prominent groups. Lipman et al. [2010] cluster sample points taken from an input shape with respect to a symmetry correspondence matrix (SCM). In the intrinsic setting, each SCM entry measures how symmetric two points are, based on the geometric similarity (a necessity for symmetry) between the local neighborhoods of the points. Xu et al. [2009] let point pairs vote for their intrinsic reflection symmetry axes and perform symmetry grouping via region growing. However, neither approach considered scale information.




Figure 2: Examples of regions possessing intrinsic symmetries and two symmetric point pairs, $\{x, x'\}$ and $\{y, y'\}$, which support the symmetries. The distance-based symmetry criterion dictates that the geodesics shown in the same color have the same length.

Our key idea for multi-scale symmetry detection is to decouple scale extraction and symmetry extraction. To this end, we perform two levels of clustering with the first identifying symmetry scales and the second extracting partial symmetries per scale. The scale clustering step is performed on point *pairs* sampled from the input shape, where point pairs which belong to intrinsic symmetries at similar scales are grouped. By clustering point pairs rather than points, surface regions covered by points in different scale clusters can overlap and they generally do. In the second step, we work within each scale cluster and cluster sample points to obtain partial symmetries at the corresponding scale. The detected multi-scale symmetries are quite general, as shown in Figure 1.

Symmetric point pairs. The most fundamental question in symmetry detection is how symmetric a pair of points are. The challenge in identifying potentially symmetric point pairs is to come up with strong conditions to adequately constrain the symmetry search so that it is computationally tractable. Instead of relying on local criteria such as geometric similarity, we judge how symmetric a pair $\{a, b\}$ of points are by measuring the *global symmetry support* the pair receives. We refer to a point pair whose symmetry support is sufficiently large simply as a *symmetric point pair*. The symmetry support for $\{a, b\}$ is estimated by the number of other point pairs which potentially share the same intrinsic symmetry as $\{a, b\}$. This latter determination is made by testing a simple, distance-based symmetry criterion on two point pairs; see Figure 2 for a few examples and Section 3 for a theoretical analysis.

Overview. The input to our algorithm is either a 2D shape defined by a closed contour or a 3D shape given by a 2-manifold mesh. The intrinsic distances are geodesic distances for 3D shapes and inner distances [Ling and Jacobs 2007] for 2D shapes. The algorithm (see Figure 3) consists of two levels of spectral clustering, operating on a set of symmetric point pairs sampled from the input shape:

- *Sampling symmetric point pairs* (Section 4):

For symmetry scale clustering, we only consider symmetric point pairs. Prior to clustering, we sample a set of symmetric point pairs from the input shape. To estimate the symmetry support received by a point pair π , we perform a *randomized voting* by counting the number of randomly chosen point pairs which potentially share the same intrinsic symmetry as π .

- *Scale extraction via pair clustering* (Section 5):

We introduce the *symmetry scale matrix* (SSM) for scale analysis. The SSM encodes symmetry affinities between the symmetric point pairs obtained. Each entry of the SSM is scale-dependent — it gives the likelihood two point pairs belong to (possibly different) intrinsic symmetries at the same scale. We perform spectral clustering using the SSM which returns different symmetry scale clusters.

We define a scale-aware signature for each symmetric point

pair π , called the *intrinsic distance profile* (IDP). The IDP encodes information about the intrinsic distances between pairs of points that belong to the same symmetry as π and characterizes the scale of that symmetry. The SSM is defined by affinities between the IDPs of the point pairs.

- *Symmetry extraction via point clustering* (Section 6):

Given a set of points belonging to (point pairs in) the same symmetry scale, we perform a second-level spectral clustering to detect symmetries at that scale. We construct a *new* SCM where the symmetry affinity measure between sample points is defined by the distance-based symmetry support. The symmetry support is again estimated via voting.

Contributions. Our main contributions are:

- A global, distance-based approach to identify potentially symmetric point pairs by measuring symmetry support.
- Multi-scale symmetry detection enabled by scale clustering of sample point pairs using scale-aware symmetry affinities.
- A robust, voting-based point-to-point symmetry affinity measure which enables the detection of partial and general intrinsic symmetries via spectral clustering.

Our algorithm leads to the extraction of prominent and overlapping intrinsic symmetries at multiple scales; see Figure 1. We demonstrate results on complex shapes exhibiting rich symmetries which vary in form and scale. We also show an application of the results to hierarchical segmentation. However, as an approach that is based on measuring intrinsic distance, it can be influenced by “topological shortcuts” which could alter the distance. It is worth noting that in Figure 1, each dancer has a single foot planted on the base, i.e., no topological discrepancy.

2 Related work

The collection of papers on symmetry detection has grown substantially in recent years [Mitra et al. 2012]. We shall not be exhaustive but only discuss works more closely related to ours. Methods for global intrinsic symmetry detection include [Ovsjanikov et al. 2008] and [Kim et al. 2010], which apply specific transforms to facilitate the symmetry search. Other works take a more direct approach by explicitly searching for the maximal distance-preserving self-maps, most notably in [Bronstein et al. 2009; Raviv et al. 2010]. The search is generally expensive, and more importantly, these works do not detect nested symmetries at multiple scales.

For partial intrinsic symmetry detection, Xu et al. [2009] explicitly extract partial intrinsic reflectional symmetry axes over a shape and Ben-Chen et al. [2010] analyze the discrete Killing vector field to detect local cylindrical symmetries. Our method is able to detect other forms of symmetries. The distance-based voting scheme in our paper draws inspiration from the work of Xu et al. [2009]. However, we generalize it for more general intrinsic symmetry detection and adapt it to multi-scale symmetry analysis.

We are not aware of existing works on multi-scale intrinsic symmetry detection. Wang et al. [2011] extract hierarchical extrinsic symmetries for man-made objects. Their analysis is at the part level and seeks a compact structural representation of a given shape. The symmetries associated with the nodes of the hierarchy do not overlap. Simari et al. [2006] introduce the folding mesh hierarchy, which recursively divides a given shape along the dominant reflectional symmetry axis. Our work deals with the more difficult problem of detecting more general intrinsic symmetries; the detected multi-scale symmetries can be processed to produce a hierarchy.




Figure 3: Outline of our multi-scale symmetry detection algorithm. After a voting step which identifies a set of (sufficiently) symmetric sample point pairs (a), we perform clustering of these point pairs based on a scale-aware affinity matrix (the SSM) to determine scale clusters. In each scale cluster, we perform the second-level clustering of sample points to detect symmetries at that scale (c-d).




Figure 4: Challenges of partial and multi-scale symmetry analysis. (a) Local geometric similarity can be unreliable as a necessary condition for symmetry: two symmetric points (red) may have dissimilar local signatures. (b) A single clustering based on point-wise symmetry correspondence [Lipman et al. 2010] may fuse different symmetries and return false positives: the three dancers (each symmetric on its own) are deemed to possess an intrinsic symmetry, yet the set does not possess an isometric intrinsic self-map.

Transformation-space voting for symmetry detection [Mitra et al. 2006; Pauly et al. 2008] is able to extract overlapping symmetries. However, this is only applicable in the extrinsic setting since only in this setting is the dimensionality of the transformation space tractable. One may be tempted to utilize the spectral transform to turn intrinsic symmetries into extrinsic ones, as done by Ovsjanikov et al. [2008] for global symmetry detection, and then apply a transformation-space voting. Unfortunately, partial intrinsic symmetries generally do not turn into extrinsic ones in the spectral embedding space. The voting approach by Podolak et al. [2006] can also detect overlapping symmetries, however it is only applicable to extrinsic reflectional symmetries.

Partial symmetry detection is related to partial matching. Many works have been proposed for partial matching, some presented in the context of symmetry detection, e.g., [Pauly et al. 2008; Bokeloh et al. 2009; Berner et al. 2011], and some not, e.g., [Gal and Cohen-Or 2006]. These methods all work in the extrinsic setting and via localized feature analysis; they are designed to extract partial matchings at small scales. Intrinsic approaches such as [Bronstein et al. 2009; Raviv et al. 2010] can detect partial matchings between large sub-shapes under articulation, but the results cover maximal regions and are not multi-scale. Our definition of partial intrinsic symmetry can be seen as a restricted case of the most general form of partial matching (see Section 3). By restricting the definition of the partial intrinsic symmetries we seek, we are able to achieve multi-scale analysis results and more discrimination.

Most closely related to our approach is the work by Lipman et

al. [2010], which also relies on spectral clustering. The top eigenvectors of their geometric similarity based SCM characterize orbits, where each orbit includes all points symmetric with one another. However, their work is not suited for multi-scale partial symmetry detection. First, exploiting local point similarities as symmetry invariants is only appropriate for global intrinsic symmetry detection. In the partial setting, it is not always reliable to judge whether two points are symmetric by comparing their point signatures, e.g., when one point lies on the boundary of symmetric regions, as shown in Figure 4(a). Second, a single clustering step is unable to identify overlapping symmetries that a point belongs to. In our work, we perform two levels of clustering.

Finally, clustering sample points using the SCM in their work is not scale-aware since the symmetry affinity in their SCM is based on comparing local shape signatures and does not reveal scale information: two symmetric points can be close or far apart. Hence, it is possible that sample points belonging to multiple symmetries are clustered together, e.g., see Figure 4(b) and contrast it with our multi-scale results in Figure 3(d). In our work, both the pair-to-pair symmetry affinity in the SSM and the point-to-point symmetry affinity in our new SCM are scale-dependent.

3 Theoretical setup

In this section, we give a more formal description of the problem we address. In particular, we define the input, the output, and the kind of intrinsic symmetries our algorithm is designed to detect — they are *isometric involutions*. We derive a distance-based symmetry criterion, a *provable necessary condition* for isometric involutions, which plays an important role in our symmetry analysis.

The reason for restricting the focus to isometric involutions is two-fold. First, having a provable necessary condition is of theoretical interest. Second, the symmetry criterion bounds the symmetry search, making it computationally tractable. That being said, we observe that in the context of partial symmetries, the restriction to involutions is not nearly as limiting as might first be suspected. We discuss this further at the end of this section and in Section 8.

Input shapes. Our problem domain is a compact, connected 2-manifold, \mathcal{M} , with or without boundary. \mathcal{M} is either a “2D shape”, i.e., $\mathcal{M} \subset \mathbb{R}^2$ is a planar region enclosed by a closed contour, or a “3D shape”, i.e., $\mathcal{M} \subset \mathbb{R}^3$. In both cases, distances on \mathcal{M} are given by the intrinsic or geodesic distance, $d_{\mathcal{M}}$, where $d_{\mathcal{M}}(x, y)$ is the length of the shortest curve in \mathcal{M} which connects x and y . In the remainder of the paper, we use the term intrinsic distance and geodesic distance interchangeably.




Figure 5: Two scenarios where the symmetry criterion (1) fails to hold: $d_{\mathcal{M}}(x, y') \neq d_{\mathcal{M}}(y, x')$. In both cases, x is mapped to x' and y to y' . (a) There is a “rotational” symmetry over the region of the bumpy triangle, but the symmetry is not an involution. (b) A repeated pattern of two arrows without an intrinsic symmetry.

Output. The output of our algorithm consists of clusters of sample points on \mathcal{M} , after the two levels of clustering analysis. Inherent to the clustering approach, the output is not a complete description of the partial intrinsic symmetries, but rather just the domains (subsets $D \subset \mathcal{M}$) associated with the partial symmetries.

Symmetry as isometric involutions. The isometric automorphisms $f : \mathcal{M} \rightarrow \mathcal{M}$ form a group \mathcal{G} that acts on \mathcal{M} . In general, any subgroup \mathcal{G}' of \mathcal{G} is considered to be an *intrinsic symmetry* of \mathcal{M} . A *partial* intrinsic symmetry of \mathcal{M} is a symmetry of a closed subset $D \subset \mathcal{M}$, where distances are measured by $d_{\mathcal{M}}$. The smallest possible non-trivial subgroup \mathcal{G}' consists of two elements: $\mathcal{G}' = \{e, f\}$, where e corresponds to the identity transformation of \mathcal{M} . In this case, f is an *involution*, i.e., f is its own inverse, and we may identify the symmetry \mathcal{G}' with f itself. Involutions play a central role in our analysis. The symmetry criterion we employ to test whether two point pairs share a common intrinsic symmetry is derived from a consideration of isometric involutions defined over *geodesically convex domains*. We say that a subset $D \subset \mathcal{M}$ is geodesically convex if for any $x, y \in D$ there is a minimal geodesic between x and y which is contained in D .

By focusing on isometric involutions over geodesically convex domains, we can derive a provable necessary condition for such symmetries (see Appendix for a proof):

Lemma 1 *Let D be a geodesically convex set in a smooth surface $\mathcal{M} \subset \mathbb{R}^3$, and let $f : D \rightarrow D$ be an isometric involution. Then, for any two points $x, y \in D$, $d_{\mathcal{M}}(x, f(y)) = d_{\mathcal{M}}(y, f(x))$.*

Symmetry criterion. Lemma 1 implies a simple-to-check criterion to test whether two point pairs potentially share the same isometric involution. Specifically, given two point pairs $\{x, x'\}$ and $\{y, y'\}$ on \mathcal{M} , the distance-based criterion

$$d_{\mathcal{M}}(x, y) = d_{\mathcal{M}}(x', y') \text{ AND } d_{\mathcal{M}}(x, y') = d_{\mathcal{M}}(y, x') \quad (1)$$

is necessarily satisfied if the pairs correspond to the same isometric involution. Figure 2 shows three examples of isometric involutions that are covered by (1). This will be our distance-based symmetry criterion, which we employ both for symmetry support voting (Section 4) and the construction of the SCM (Section 6).

Undetected symmetries. Our algorithm is not designed to detect symmetries for which (1) fails to hold. Figure 5 illustrates two scenarios. In 5(a), the region of the bumpy triangle has a 3-fold rotational symmetry over a geodesically convex domain, yet (1) does not hold on the pairs $\{x, x'\}$ and $\{y, y'\}$. In 5(b), the two arrows form what may be perceived as a translational symmetry, though

they do not possess an intrinsic symmetry. The criterion (1) fails to hold in this situation. Generally speaking, our algorithm is not designed to detect all forms of *repeated patterns*, especially if the patterns are “non-symmetrically connected”. Figure 4(b) is such an example: the three dancers repeat, but the distances between them are unequal. Figure 2(c) is the opposite: the two L ’s are symmetrically connected and we can detect the involution.

Generality of symmetry criterion. Applying the symmetry criterion (1) allows us to detect more general symmetries beyond isometric involutions, as evidenced by the numerous results we obtained, e.g., see Figures 10-13. This can be attributed to a “symmetry fusion”, discussed further in Section 8. For example, any permutation can be decomposed as a product of transpositions [Clark 1984, §80]. A transposition is a transformation that interchanges two objects: it is an involution. The many symmetries we are able to identify over sets of isometric figures can be explained by a fusion of such pairwise transpositions.

Weak partial intrinsic symmetry. We finally remark on a relaxation of our partial intrinsic symmetry definition. While our definition is a commonly used one [Xu et al. 2009; Raviv et al. 2010; Mitra et al. 2012], a subtlety that arises in the context of partial symmetries should be mentioned. A symmetry of $D \subset \mathcal{M}$ with respect to the intrinsic metric of D itself, need *not* be a partial intrinsic symmetry of \mathcal{M} by our definition, since in general, for $x, y \in D$, the shortest path between x and y may not lie entirely in D . In this latter case, one can define $d_D(x, y) = \infty$ if x and y lie in different connected components of D . We call a symmetry of D with respect to d_D a *weak partial intrinsic symmetry*. Figure 5(b) and the isolated humans in Figures 4(b) give examples of weak partial intrinsic symmetries. Such a relaxed definition corresponds to partial matching, which is generally a more challenging problem.

4 Symmetric point pairs by voting

Our symmetry analysis algorithm operates on a set \mathcal{X} of n sample points, uniformly taken from a given input shape. Prior to symmetry scale clustering, which we describe in the next section, we extract a set of symmetric point pairs over \mathcal{X} ; only these pairs are clustered during scale extraction. We now describe a voting approach based on our symmetry criterion (1) to accumulate symmetry support and extract the symmetric point pairs.

Voting for one pair. Recall that a point pair π is symmetric if it has a sufficiently large symmetry support, which is measured by the number of point pairs which satisfy the criterion (1) with π . A straightforward approach is to randomly choose a pair $\alpha = \{x, x'\} \subset \mathcal{X}$ and let a randomly chosen set of sample pairs vote whether they satisfy (1) with α (a positive vote) or not (a nega-




Figure 6: Comparison of symmetry pairs voting for one pair (c) vs. for two pairs (d). A point pair $\{x, x'\}$ can belong to two distinct symmetries (a-b), in which case voting for just the pair $\{x, x'\}$ can result in pairs supporting multiple symmetries (c). Two pairs $\{x, x'\}$ and $\{z, z'\}$ have a better chance of pinning down one symmetry (d).

tive vote). A positive vote results if each distance difference is less than ϵ_d of the maximum of the two compared distances, where ϵ_d is a user-specified threshold. Note however that one point pair may belong to distinct symmetries; see Figures 6(a) and (b). We find that voting for one pair often yields many point pairs which belong to different symmetries; see Figure 6(d). This is undesirable since the symmetry support we seek is meant to correspond to the same symmetry.

Voting for two pairs. To address the above issue, we employ a voting scheme where the strategy is to vote for *two* pairs, instead of one. Specifically, we randomly sample two point pairs α and β which satisfy the symmetry criterion (1). Then for any other sampled point pair π , we determine whether it satisfies (1) with *both* α and β . The key motivation for voting for two pairs is that other pairs which cast positive votes with respect to both pairs are more likely to share a common symmetry compared to the case of voting for one pair; see Figure 6 for an illustrative example.

Let Π denote the set of all point pairs over \mathcal{X} and $N = \binom{n}{2} = |\Pi|$. If the number of point pairs casting positive votes with respect to both α and β is greater than some fraction p_e of N , then we regard this set of point pairs (including α and β) to potentially form a symmetry. We denote this set of pairs by $\Sigma^{\alpha,\beta}$ and designate the tuple $\langle \alpha, \beta \rangle$ as the *representative* of this symmetry. We repeat the above voting process with new random pairs α and β and collect all the resulting $\Sigma^{\alpha,\beta}$ into a single set $\Phi = \cup \Sigma^{\alpha,\beta}$, which provides the set of symmetry point pairs for scale clustering. The voting process stops when $|\Phi|$ reaches a fraction p_t of N .

The sets $\Sigma^{\alpha,\beta}$ are only potential symmetries, which we utilize for scale and symmetry detection; the final set of symmetries are the result of the two levels of clustering. All the sets $\Sigma^{\alpha,\beta}$ and their representatives are pre-computed and stored for the subsequent analysis steps. For any symmetric point pair $\pi \in \Sigma^{\alpha,\beta}$, we refer to $\Sigma^{\alpha,\beta}$ as its *supporting set*. If π belongs to more than one such set, we choose the largest one. Accordingly, the point pairs in $\Sigma^{\alpha,\beta}$ are referred to as the *supporting pairs* for π .

5 Multi-scale symmetry analysis

In order to sort out the partial intrinsic symmetries at different scales, we opt to cluster the symmetric point *pairs* rather than the sampled points. We define for each symmetric point pair a scale-aware signature to encode scale information. Based on the scale-aware signatures, we perform clustering over the point pairs, where each cluster contains symmetries of a particular scale.

For an intrinsic symmetry, we define its symmetry scale as that of the isometric self-map, measured by a scale or distance profile for the point pairs within the symmetry support set. Since the symmetries of different scales are unknown, computing their support region is intractable. Instead, we measure symmetry scale locally at every symmetric point pair by an encoding of intrinsic distances between pairs of points that support the symmetric point pair.

Intrinsic distance profile (IDP). Given a symmetric point pair $\{p, q\}$, we compute a signature $\mu_{\{p,q\}}$, called the intrinsic distance profile (IDP), where the distances are measured between p and q and between points in each supporting pair for $\{p, q\}$. The IDP is scale-aware since it encodes information about the intrinsic distances and scale is characterized by these distances.

To define the profile signature for $\{p, q\}$, we form l bins, $\{b_i\}_{i=1}^l$, each containing a subset of supporting point pairs for $\{p, q\}$. Distribution of the supporting pairs into the bins is based on a *pair-to-pair* distance measured from $\{p, q\}$. Specifically, the pair-to-pair

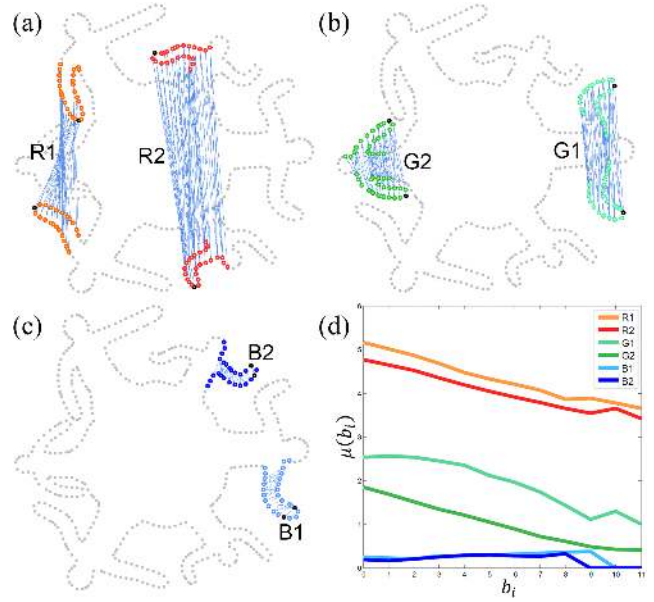


Figure 7: Examples of symmetric point pairs (black dots), their supporting pairs (linked by blue lines), and plots of IDPs. (a-c) Six symmetric point pairs and their supporting pairs whose weight contributing to the IDP is larger than 10^{-6} . (d) The corresponding IDPs (by matching colors) show their discriminative power.

distance between $\{p, q\}$ and another point pair $\{s, t\}$ is defined as

$$d_{\mathcal{M}}(\{p, q\}, \{s, t\}) = \min \left\{ \frac{d_{\mathcal{M}}(p, s) + d_{\mathcal{M}}(q, t)}{2}, \frac{d_{\mathcal{M}}(p, t) + d_{\mathcal{M}}(q, s)}{2} \right\},$$

where \min is used to account for the permutation of point labels.

Each of the l bins contains the subset of supporting point pairs for $\{p, q\}$ whose distance to $\{p, q\}$ falls within a sub-interval of $[0, L_{\mathcal{X}}]$, where $L_{\mathcal{X}}$ is the geodesic diameter of the set of sample points \mathcal{X} , $L_{\mathcal{X}} = \max_{a,b \in \mathcal{X}} d_{\mathcal{M}}(a, b)$, and hence it is the maximum possible pair-to-pair distance from $\{p, q\}$. The sub-intervals are uniformly sized. Thus bin b_i contains pairs whose distance to $\{p, q\}$ is between $[(i-1)/l] \cdot L_{\mathcal{X}}$ and $(i/l) \cdot L_{\mathcal{X}}$.

Within each bin b_i , one can encode any information about the set of point pairs belonging to that bin. We would like an encoding that is simple (for efficiency), scale-aware, and discriminative. To this end, we compute a weighted average of the *geodesic widths* for the point pairs belonging to b_i as follows,

$$\mu_{\{p,q\}}(b_i) = \frac{\sum_{\{s,t\} \in b_i} w_{\{s,t\}} d_{\mathcal{M}}(s, t)}{\sum_{i=1}^l \sum_{\{s,t\} \in b_i} w_{\{s,t\}}},$$

with a Gaussian weight $w_{\{s,t\}} = e^{-(d_{\mathcal{M}}(\{p,q\}, \{s,t\})/(\sigma_b L_{\mathcal{X}}))^2}$. Note that the geodesic width for a pair $\{p, q\}$ is simply the geodesic distance between p and q . We introduce the new term for convenience of presentation: It avoids overloading the word “distance”.

Since our algorithm focuses on partial symmetries, we obtain locality of the IDP signature by using the Gaussian weighting function to truncate the bins which are far away from $\{p, q\}$. The localization parameter, i.e., the Gaussian width, is chosen as $\sigma_b = 0.2$ in our implementation. The number of bins l is fixed at $l = 10$. Figure 7 demonstrates that the scale descriptor captures well the scale of the local symmetry supported by a point pair.

Scale-aware clustering of point pairs. With the scale-aware pair signature IDP, we can perform pair clustering based on spectral analysis. To this end, we build a dissimilarity matrix, called the symmetry scale matrix (SSM). The dimension of the SSM is $M \times M$, where M is the size of symmetric pair set Φ . Each entry is computed as the earth mover’s distance [Pele and Werman 2009] between the IDPs of the two corresponding point pairs. To automatically determine the number of clusters, we employ self-tuning spectral clustering [Zelnik-Manor and Perona 2004], where the eigen-decomposition of the symmetry scale matrix is computed with ARPACK [Lehoucq et al. 1998]. After clustering, the set of symmetric point pairs Φ is divided into k clusters: $\Phi = \cup\{\Phi_i\}_{i=1}^k$, where each contains symmetries of the corresponding scale. Figure 8 (b-d) demonstrates a result of pair clustering.

Alternatives to IDP. The definition of IDP requires computing and binning intrinsic distances and the signature is a vector. A simpler signature would be the average geodesic width taken over all the supporting pairs. Even simpler would be the intrinsic distance between p and q . However, compared to the vector-form IDP, the two simpler alternatives are not sufficiently descriptive of the symmetries to which they belong since they are based on single numerical values. In particular, the IDP along with the earth mover’s distance can capture the local geometry of a supporting pair set while the simpler alternatives are too coarse. Experimentally, we have found the results from these simpler signatures to be unsatisfactory.

6 Partial intrinsic symmetry extraction

In this section, we introduce the construction of the symmetry correspondence matrix (SCM), which captures the intrinsic symmetries of a particular scale based on the corresponding pair cluster.

Definition of intrinsic SCM. As in [Lipman et al. 2010], we define the dissimilarity measure for detecting intrinsic symmetries between a pair of points $x_i, x_j \in \mathcal{X}$ by considering all possible intrinsic mappings f that take x_i to x_j ,

$$S_{ij} = S(x_i, x_j) = \inf_{f \in \mathcal{F}: f(x_i)=x_j} D_f(\mathcal{X}, f(\mathcal{X})),$$

where \mathcal{F} denotes the set of intrinsic self-mappings and

$$D_f(\mathcal{X}, f(\mathcal{X})) = \left(\frac{\sum_{p,q \in \mathcal{X}} (d_{\mathcal{M}}(p, q) - d_{\mathcal{M}}(f(p), f(q)))^2}{N} \right)^{\frac{1}{2}},$$

measures the intrinsic deviation between two point sets \mathcal{X} and $f(\mathcal{X})$, with N being the number of all point pairs. To accommodate partial intrinsic self-mappings, D_f should return a low distance if f matches some sufficiently large part of \mathcal{X} . This has been reflected in the symmetric point pairs voting (Section 4) where we used the fraction p_e for identifying all potential intrinsic symmetries.

The dissimilarity matrix $S \in R^{n \times n}$, where n is the number of sample points, is then converted into a symmetry correspondence matrix (SCM) $C \in R^{n \times n}$ using a Gaussian kernel $C_{ij} = e^{-(S_{ij}/(\sigma S_{max}))^2}$, where we use $\sigma = 0.01$, and S_{max} is the maximum value of all S_{ij} .

Computation of dissimilarity matrix S . One needs to find the proper self-mapping to compute the SCM. For the extrinsic case, the self-mappings are rigid transformations each of which is found through aligning the local frames at two points and observing how much the shape is aligned by the transformation [Lipman et al.




Figure 8: Symmetry clusters (shown with different colors) detected with all symmetric point pairs do not reveal multiscale symmetries (a). (b-d) show the symmetric point pairs in three scale clusters and the corresponding multi-scale symmetries detected in each scale.

2010]. However, intrinsic mappings cannot be easily parameterized and computed this way. We instead rely on the discretized approximation of intrinsic self-mappings represented by the symmetric point pairs Φ we have computed in Section 4.

Suppose we wish to detect symmetries for the scale cluster Φ_i . We first initialize the entries of S as $S_{ij} = \infty$. For each representative (α, β) and the corresponding supporting set $\Sigma^{\alpha, \beta}$, if $\Sigma^{\alpha, \beta}$ is sufficiently covered by the current scale cluster, specifically $|\Sigma^{\alpha, \beta} \cap \Phi_i| > 0.9 |\Sigma^{\alpha, \beta}|$, we measure the intrinsic deviation for the intrinsic self-mapping $f^{\alpha, \beta}$ supported by $\Sigma^{\alpha, \beta}$ as the averaged local deviations of all point pairs in $\Sigma^{\alpha, \beta}$ w.r.t. α and β :

$$D_{f^{\alpha, \beta}} = \left(\frac{\sum_{\pi \in \Sigma^{\alpha, \beta}} \max\{\delta(\alpha, \pi), \delta(\beta, \pi)\}^2}{|\Sigma^{\alpha, \beta}|} \right)^{\frac{1}{2}}, \quad (2)$$

where $\delta(\cdot, \cdot)$ is the local deviation measured between two point pairs derived from distance criterion (1):

$$\delta(\{x, x'\}, \{y, y'\}) = \max\{|d_{\mathcal{M}}(x, y) - d_{\mathcal{M}}(x', y')|, |d_{\mathcal{M}}(x, y') - d_{\mathcal{M}}(x', y)|\}.$$

We then update the entries of S_{pq} for all $\{p, q\} \in \Sigma^{\alpha, \beta}$ as $\min\{S_{pq}, D_{f^{\alpha, \beta}}\}$.

Upper bound of entries of S . Similarly to [Lipman et al. 2010], we find an upper bound for those entries which are not updated instead of leaving them as ∞ . Looking for a good upper bound in the entire space of intrinsic mapping is by no means easy. However, it is possible to find an upper bound within the space of rigid transformations which is a subspace of the intrinsic transformations.

We therefore consider a translation $t^{i,j} = x_j - x_i$ which takes x_i to x_j and compute $D(\mathcal{X}, \mathcal{X} + t^{i,j})$ as the upper bound for S_{ij} . However, the estimation of this upper bound is not trivial as in the extrinsic case. Specifically, we first apply $t^{i,j}$ on \mathcal{X} and compute a




Figure 9: Symmetry detection on 2D shapes sorted by scale. Observe the inter- and intra-object symmetries detected at multiple scales, even down to the small scales of the limbs in (d).

point correspondence by finding the mutually closest points. Using this point correspondence, we can compute the intrinsic deviation $D_{t^i,j}$ using Equation (2). At the end of the voting process, we translate the minimal value of S to zero and take the minimum of each entry of S_{ij} with $D_{t^i,j}$. We find this upper bound is easy to compute and works well in practice.

Spectral analysis of scale-aware SCM. After constructing the SCM for a particular scale, we perform spectral analysis to detect the partial intrinsic symmetries in that scale. Here we again adopt self-tuning spectral clustering. For the 2D shape in Figure 8 (a), we demonstrate the multiscale symmetries detected via spectral analysis over the scale-aware symmetry correspondence matrices; see Figure 8 (b-d). As shown in the figures, our method robustly detects symmetries at different scales.

For a given scale, there might exist some symmetry clusters encompassing very few point pairs belonging to that scale, meaning that they are not well supported and do not reveal any symmetry of the particular scale. Such *void clusters* should be discarded in the final results. We employ a simple scheme to identify the void clusters: for the clusters obtained in a given scale, we measure the voidness of a symmetry cluster by computing its point pair coverage, based on all point pairs belonging to the scale. Specifically, if the pair coverage of a cluster is less than 10% of the sample points in the whole scale cluster, we regard it as a void cluster. In figures throughout the paper, void clusters are shown in gray in the 2D cases and in transparent rendering in the 3D cases.

7 Results

In this section, we demonstrate and discuss results obtained by our multi-scale partial intrinsic symmetry detection algorithm on both 2D and 3D shapes. Comparison to closely related approaches [Xu et al. 2009; Lipman et al. 2010] is also given. We also show how the detected symmetries can be utilized to obtain hierarchical segmentations that better reveal the semantics of a shape. Additional results can be found in the supplementary material.

All the experimental results shown in the paper were obtained by using the same set of parameters. Specifically, the threshold used for testing the criterion (1) is $\epsilon_d = 0.03$. The fractions used in symmetric pairs voting are: $p_e = 0.003\%$ and $p_t = 0.06\% \sim 0.1\%$ (depending on the smallest scale the user wishes to detect).

Pre-processing. Geodesic distances for 3D models are computed using the algorithm and available implementation of Surazhsky et al. [2005]. Some 3D models contain close-by but disconnected parts. Since we require the input mesh to be connected, we

fuse the disconnected parts. In cases where a model contains too many disconnected, we resort to voxelization and use inner distance in the constructed volume as the intrinsic distance for symmetry detection. For sample point selection, we use the uniform sampling module from MeshLab.

Segmentation enhancement. A natural consequence of using a sampling and clustering approach is that the cluster boundaries may be imperfect. Also, the symmetries are sometimes imperfect and only approximate. To obtain higher-quality boundaries, we apply a post-processing step which we call segmentation enhancement. We first over-segment the models using any shape segmentation method based on the minima rule [Shamir 2008] and then merge the segments with the corresponding symmetry clusters we have detected in various scales. All the detection results on 3D models shown in the paper were obtained using segmentation enhancement unless otherwise noted (see Figures 10, 11 and 12).

In most cases, the raw clustering results are very close to being identical to those from segmentation enhancement, except for the case of imperfect symmetries, as shown in Figure 12. We observe that the symmetries covering the small fingers all have correct boundaries. The imperfect boundaries in (a) and (b) can be attributed to the size discrepancy between the arms.

Symmetry detection. Figures 3, 9, and 18 demonstrate several 2D results and Figures 10-13 show results for 3D models with intricate structures. In each displayed model, a sub-shape shown in the same color represents a detected intrinsic symmetry. The gray color (in 2D) and transparent rendering (in 3D) indicate regions which are not covered by the current scale cluster. There are several observations to be made from these results:

- *Multi-scale symmetries:* The ability of our method to identify multiple and nested partial intrinsic symmetries is evident. The extracted symmetry scales range from those of groups of objects undergoing approximately isometric deformations to those that cover individual components of an object that possess various forms of self-symmetry, e.g., individual limbs in 2D having reflection-like symmetry (see Figure 9), and similarly for 3D limbs (see Figures 10-14).
- *Intra- and inter-object symmetries:* Our method is able to detect the self-symmetry in an integral object such as a human, as well as inter-object symmetries. The former can lead to meaningful object identification while the latter allows one to recognize groups of similar objects.
- *General symmetries:* Observe that the arms of the octopus in Figure 10 and the thumbs of the hand-arm model in Figure 12 represent symmetry groups that are more general than those detected by the algorithm of Xu et al. [2009], which analyzes intrinsic reflectional symmetries only. For example, one cannot find a proper reflectional symmetry axis on the surface of the hand-arm model that links the two thumbs.
- *Disconnected symmetries:* The symmetry regions detected are not necessarily geodesically convex; they may be disconnected, e.g., see several examples of distinct figures being detected in Figures 10-13. In these cases, two distinct figures may be connected by a thin bridge of geodesics joining them, thus forming a convex domain. However, the bridge between them may be too thin to be captured by the sampling resolution. Hence only the disconnected regions are returned.
- *Noise:* Our method is statistical in nature, with the use of randomized voting and distance averages for IDP, for example. Hence it is fairly robust against reasonable levels of noise. Figure 11 demonstrates such robustness on two models with synthetic Gaussian noise added.



Figure 10: A gallery of 3D symmetry detection results sorted by scale. From left to right and top-down: Children, Octopus, Kung Fu Panda, IndoLady, and Thai Statue. The last two models show raw clustering results. Observe again the inter- and intra-object symmetries detected at a significant range of scales. The IndoLady and Thai Statue models were chosen to demonstrate the performance of our algorithm on models which are not compositions of articulated characters. For the children model, not all arms or legs are detected in the last shown scale due to scale discrepancies and some parts fused with the body. For the IndoLady, our method does not return all perceived symmetries, e.g., the self-symmetries of the individual limbs. To save space, the first image for the Kung Fu Panda contains the first five scales, each revealing a self-symmetry of the four characters and the base.

Statistics. Our experiments were performed on an Intel Core (TM) 3.40GHz machine with 4.00GB RAM. For all the 2D examples, we employ 500 samples. For 3D models, the number of sample points is about 3K with small discrepancies due to criteria employed by the MeshLab sampling routine. Figure 13 shows that more samples allow finer-scale symmetries to be detected but at a cost of processing times. Table 1 reports various statistics including timing. The most time consuming part is the SSM analysis, of which at least 85% of the time was spent on the computation of the earth mover’s distance for the profile signatures (IDP).

Comparisons. The two methods most closely related to ours are [Xu et al. 2009] and [Lipman et al. 2010]. The former can detect partial intrinsic reflectional symmetries while the later is designed



Figure 11: Raw symmetry clusters in multiple scales detected on a noisy Neptune model and a noisy IndoLady model. The noise level is 1.0, which is the ratio of the average vertex displacement over the average edge length in the original mesh.




Figure 12: Raw clustering results for symmetry detection. Imperfect boundaries (a-b) are due to discrepancy between the arms.

to deal with partial extrinsic symmetries. However, if the symmetric sub-shapes do not undergo significant pose variations, the global alignment component of [Lipman et al. 2010] may allow it to detect certain partial intrinsic symmetries. Regardless of the case, both methods can only detect partial intrinsic symmetries that only provide a *single* coverage of the input shape. Our method, on the other hand, detects nested intrinsic symmetries which include symmetries detected by these two methods and more, as shown in Figure 14.




Figure 13: Denser sampling results in finer-scale symmetries detected. (a-c) 2,016 sample points. (d) 3,080 sample points.

Model	#s	t_{vote}	#sp	t_{ssm}	#l	t_{scm}
Ballet (1)	3,260	59	6,045	915	5	571
Family (9)	500	7	2,294	234	4	16
Birds (9)	500	5	1,760	189	4	13
Children (10)	2,933	50	5,845	852	4	417
Octopus (10)	3,045	58	6,175	939	2	273
Panda (10)	2,817	42	5,906	929	4	442
IndoLady (10)	2,227	36	4,209	594	2	174
Thai (10)	3,136	43	5,268	767	3	343
Neptune (11)	2,742	39	5,848	874	4	361
Arm (12)	2,563	38	5,441	763	4	381
Soldier (13)	2,016	32	3,923	371	3	216
Soldier (13)	3,080	42	5,657	816	4	402
Dancer (16)	2,600	46	5,802	836	4	349

Table 1: Statistics on some experiments. The numbers appearing in parentheses refer to figure numbers in the paper. #s: number of sample points; #sp: number of symmetric point pairs; #l: number of scales detected. Running times are reported in seconds for symmetric pairs voting (t_{vote}), SSM, and SCM analyses; they do not account for intrinsic distance computation.



Figure 14: Comparison between our multi-scale results (a) to the symmetry detection results of [Lipman et al. 2010] (b) and [Xu et al. 2009] (c), which both provide only a single coverage of the shape. Our method detects overlapping symmetries and these results combine those from the other two methods.

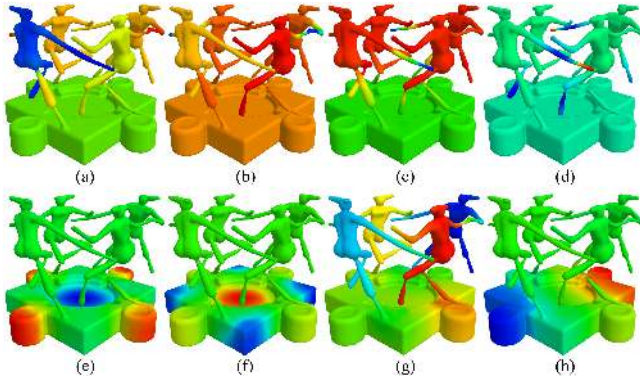


Figure 15: Plots of the top two eigenvectors of GPS (a-b), HKS (c-d), and our multiscale partial intrinsic SCM in two scales (e-h). The partial intrinsic symmetries are more clearly revealed in (e-h).

Symmetry invariants, e.g., the Global Point Signature (GPS) [Rustamov 2007] and the Heat Kernel Signature (HKS) [Sun et al. 2009]), have been used for global intrinsic symmetry detection [Lipman et al. 2010]. However, they are not effective in the partial intrinsic setting. Figure 15 compares the top two eigenvectors of our multiscale partial intrinsic SCM (in two scales) against those of the GPS and HKS. The plot shows that our method can more clearly reveal the partial intrinsic symmetries at different scales.

Hierarchical segmentation. As several previous works [Simari et al. 2006; Podolak et al. 2006; Xu et al. 2009] have demonstrated, symmetry, as a high-level cue, can lead to more semantic segmentation results since common objects all possess self-symmetries. With our multi-scale approach, we can easily obtain a symmetry-driven hierarchical segmentation using the following scheme. We first construct a directed graph whose nodes are the segments (including the whole shape) induced by the multiscale symmetries and whose directed edges encode the spatial inclusion relations between the segments. Specifically, there is an in-edge from node A to B if segment A is a subset of segment B . Then starting from the node representing the whole shape, we repeatedly remove the nodes (one at a time) with only zero in-degree and all its associated out-edges. The sequence of removals reflects the levels of segments, which is then used to construct the hierarchy in a top-down manner. Figure 16 shows such a hierarchy for the Dancer model.

Figure 17 shows comparison results between our segmentation and those from two state-of-the-art algorithms. The first is the hierarchical segmentation based on fitting primitives [Attene et al. 2006]. The second is the randomized cut approach of [Golovinskiy and Funkhouser 2008], which is not hierarchical. We show results ob-




Figure 16: A hierarchical organization of detected symmetries on the Dancer model. The hierarchy can be used to derive a hierarchical segmentation of the model (see Figure 17).




Figure 17: Comparison between our symmetry-driven hierarchical segmentation scheme (top row) and hierarchical segmentation based on primitive fitting [Attene et al. 2006] (middle row) and normalized cut [Golovinskiy and Funkhouser 2008] (bottom row). Each column shows the same segmentation count. It is evident that our results conform better to the shape semantics.

tained for the same segmentation counts. As can be seen, with symmetry considerations, our results are qualitatively superior, beyond boundary qualities, with better conformation to shape semantics.

It is interesting to consider the usefulness of our segmentation results for object identification and partial matching in the large scale. Past works on repeated pattern detection using symmetry analysis [Mitra et al. 2006; Pauly et al. 2008] are designed for such purposes. They typically rely on localized analysis and are able to detect repeated patterns at the low level. The same can be said about approaches developed for partial matching, e.g., [Gal and Cohen-Or 2006], as they often cannot return a partial matching between large-scale objects. We believe that our approach is able to accomplish this task due to its multi-scale nature.

Limitation. Perhaps the most fundamental limitation of our approach should be attributed to its reliance on global intrinsic distances. In the presence of a topological “shortcut” (see Figure 18(a) for an example), perhaps due to noise, perceived symmetric regions may have topological discrepancies. As a result, certain inter-object




Figure 18: Limitation to the use of intrinsic distances (a 2D case). Having one foot of the right figure planted into the base (a) or disconnected from the base (b) has a drastic effect on the multi-scale symmetries detected, since the distances changed drastically.



Figure 19: Limitation to the use of intrinsic distances (a 3D case). Our method returns erroneous symmetry regions in the first two scales (a-b). The input model is the same as the one in Figure 1 except that the last dancer now has both feet planted into the base, creating a topological inconsistency among the three dancers. Symmetry detected for the remaining scales are the same as Figure 1.

symmetries may not be detected as we might expect. Figure 19 shows a similar case in 3D, where comparing to Figure 1, now the last dancer has *both* feet planted into the base, making it topologically inconsistent with the other two dancers. The cases involving topological shortcuts are encompassed by weak partial intrinsic symmetries: the two outermost dancers in Figure 19 represent such a symmetry. Sensitivity to topological changes is a fundamental limitation of any approach that relies on global intrinsic distances for shape analysis. With the presence of topological noise, the use of defect-tolerant geodesics [Campen and Kobbelt 2011] may provide a remedy.

Our method is purely geometry-based without considering semantic knowledge, hence some perceived symmetries may not be detected. Also limited by the reliance on sampling and clustering, we cannot hope to detect all possible symmetry scales. For example, in the IndoLady model in Figure 10, we do not return the self-symmetries of the individual limbs. In the children model, in the last scale, one pair of legs (of the first child) is missed. Additional limitations of our current approach are discussed in the next section.

8 Conclusion, discussion, and future work

We present an algorithm for multi-scale partial intrinsic symmetry detection on 2D and 3D shapes. While the results obtained are quite encouraging, we still regard our work as only a preliminary attempt at a complete and verifiable solution to the general problem. We now discuss additional issues pointing to limitations of our approach and possible future works to address these issues.

General symmetries beyond isometric involutions. Although our symmetry criterion is designed for detecting isometric involutions over geodesically convex domains, in practice we are able to

detect symmetries that do not correspond to involutions or whose domain is not convex. The detected symmetry encompassing three soldiers in Figure 13 provides an example. An explanation is that pairwise transpositions transitively create a high correlation between the three figures, whereas the bridges of geodesics joining the figures in pairs are not correlated together. Therefore to some extent, our algorithm can capture weak partial intrinsic symmetries. This demonstrates an effect of “symmetry fusion”, the fusion of the pairwise transpositions, which is the result of applying clustering.

Symmetry fusion. Multi-scale analysis using pair clustering does not completely prevent symmetry fusion from occurring in the SCM clustering step. When multiple symmetries share the same scale, implying that their sample points would belong to the same scale cluster, SCM analysis may fuse these symmetries together. An example is the children model shown in Figure 10. Symmetries between the adjacent children all share the same scale (the third scale in the children sequence). The algorithm returns the fused point cluster but not the children pairs.

From clusters to self-maps. One of the inherent limitations of clustering-based symmetry analysis is that the resulting symmetry is represented by a set (of samples) and it is typically not associated with an isometric automorphism (the symmetry map). On the other hand, extrinsic approaches relying on partial matching or transformation-space voting do return the symmetry maps. However, the computation of the symmetry maps is made easy by confining the search to a given region. Once the self-symmetric region is known, we are left with the problem of global intrinsic symmetry detection to which existing methods such as [Raviv et al. 2010; Ovsjanikov et al. 2008] can potentially be applied to obtain the maps.

The notion of scale. The scale of a symmetry is perhaps more intuitively equated to its size or area of coverage. Our scale analysis however is not based on area, but intrinsic distances. This is mainly attributed to the fact that our definition of partial intrinsic symmetry is distance-based and so is our detection scheme. There is no direct correlation between the area of a symmetry coverage and distances between symmetric points. For example, two symmetric regions that are very small in area but far apart geodesically are considered to possess a large scale by our SSM analysis. Characterizing and clustering symmetries based on their area of coverage appears to be more challenging. With incomplete information about the symmetries, it is difficult to reliably estimate the area of a region from a sampling of points or point pairs.

Small-scale symmetries. When a symmetry covers a very small region, there may not be sufficient sampling over that region. The clustering scheme may then be unable to detect such a small cluster. That said, symmetry detection schemes are generally more focused on identifying prominent symmetries rather than ones that are very local. For example, any geodesic circular region possesses an intrinsic symmetry, however small the region is, but we are often not interested in returning such small-scale symmetries.

Future work. Besides addressing the above issues, we would like to improve the performance of our algorithm. Parallelization is always a possibility since the entries in the SSM or SCM can all be computed in parallel. It would be desirable to enhance the algorithm so as to detect weak partial intrinsic symmetries. In this case, a purely distance-based approach may need to be augmented with additional considerations. Finally, it would be useful to extend our method to work on imperfect and even incomplete geometry data.

Acknowledgements. We would first like to thank the anonymous reviewers for their valuable feedback. Thanks also go to Daniel Cohen-Or for fruitful discussions on the paper. Part of the 3D models in this paper is from the shape repositories of AIM@SHAPE and Stanford. This work is supported in part by grants from NSFC (61202333, 61232011, 61161160567, 61025012, 61103084, and 61070071), NSERC (No. 611370), National 863 Program (2011AA010503), Shenzhen Science and Innovation Program (CXB201104220029A, JC201005270329A), the 973 National Basic Research Program of China (2011CB302400).

Appendix: Proof of Lemma 1.

Proof Let $\alpha : [0, \ell] \rightarrow D$ be an arclength parameterized minimizing geodesic between x and $f(y)$. Thus $\alpha(0) = x$, $\alpha(\ell) = f(y)$, and $\ell = \ell(\alpha) = d_{\mathcal{M}}(x, f(y))$, where $\ell(\alpha)$ denotes the length of α . Now consider the curve $\tilde{\alpha} = f \circ \alpha$. Since f is an isometry, $\tilde{\alpha}$ is a geodesic of the same length as α . Also, we have $\tilde{\alpha}(0) = f(x)$, and $\tilde{\alpha}(\ell) = f(f(y)) = y$, where the last equality follows from the assumption that f is an involution.

Thus $d_{\mathcal{M}}(f(x), y) \leq \ell(\tilde{\alpha}) = d_{\mathcal{M}}(x, f(y))$. A symmetric argument starting with a minimizing geodesic between y and $f(x)$ yields the reverse inequality, and the result follows. \square

References

- ATTENE, M., FALCIDIENO, B., AND SPAGNUOLO, M. 2006. Hierarchical mesh segmentation based on fitting primitives. *The Visual Computer* 22, 3, 181–193.
- BEN-CHEN, M., BUTSCHER, A., SOLOMON, J., AND GUIBAS, L. 2010. On discrete Killing vector fields and patterns on surfaces. *Computer Graphics Forum (Special Issue of SGP)* 29, 5, 1701–1711.
- BERNER, A., WAND, M., MITRA, N. J., MEWES, D., AND SEIDEL, H.-P. 2011. Shape analysis with subspace symmetries. *Computer Graphics Forum (Special Issue of Eurographics)* 30, 2, 277–286.
- BOKELOH, M., BERNER, A., WAND, M., SEIDEL, H.-P., AND SCHILLING, A. 2009. Symmetry detection using line features. *Computer Graphics Forum (Special Issue of Eurographics)* 28, 2, 697–706.
- BRONSTEIN, A. M., BRONSTEIN, M. M., BRUCKSTEIN, A. M., AND KIMMEL, R. 2009. Partial similarity of objects, or how to compare a centaur to a horse. *Int. J. Comp. Vis.* 84, 2, 163–183.
- CAMPEN, M., AND KOBELT, L. 2011. Walking on broken mesh: Defect-tolerant geodesic distances and parameterizations. *Computer Graphics Forum (Special Issue of Eurographics)* 30, 2, 623–632.
- CLARK, A. 1984. *Elements of Abstract Algebra*. Dover, New York.
- GAL, R., AND COHEN-OR, D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Trans. on Graph* 25, 1, 130–150.
- GOLOVINSKIY, A., AND FUNKHOUSER, T. 2008. Randomized cuts for 3d mesh analysis. *ACM Trans. on Graph* 27, 5, 145:1–12.
- KIM, V. G., LIPMAN, Y., CHEN, X., AND FUNKHOUSER, T. 2010. Möbius transformations for global intrinsic symmetry analysis. *Computer Graphics Forum (Special Issue of SGP)* 29, 5, 1689–1700.
- LEHOUCQ, R., SORENSEN, D., AND YANG, C. 1998. *ARPACK users' guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods*, vol. 6. SIAM.
- LING, H., AND JACOBS, D. W. 2007. Shape classification using the inner-distance. *IEEE Trans. Pat. Ana. & Mach. Int.* 29, 286–299.
- LIPMAN, Y., CHEN, X., DAUBECHIES, I., AND FUNKHOUSER, T. 2010. Symmetry factored embedding and distance. *ACM Trans. on Graph* 29, 4, 103:1–12.
- MITRA, N. J., GUIBAS, L. J., AND PAULY, M. 2006. Partial and approximate symmetry detection for 3D geometry. *ACM Trans. on Graph* 25, 3, 560–568.
- MITRA, N. J., PAULY, M., WAND, M., AND CEYLAN, D. 2012. Symmetry in 3D geometry: Extraction and applications. In *Proc. of Eurographics STAR Report*.
- OVSJANIKOV, M., SUN, J., AND GUIBAS, L. 2008. Global intrinsic symmetries of shapes. *Computer Graphics Forum (Special Issue of SGP)* 27, 5, 1341–1348.
- PAULY, M., MITRA, N. J., WALLNER, J., POTTMANN, H., AND GUIBAS, L. 2008. Discovering structural regularity in 3D geometry. *ACM Trans. on Graph* 27, 3, 43:1–11.
- PELE, O., AND WERMAN, M. 2009. Fast and robust earth mover's distances. In *Proc. Int. Conf. on Comp. Vis.*, 460–467.
- PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A planar-reflective symmetry transform for 3D shapes. *ACM Trans. on Graph* 25, 3, 549–559.
- RAVIV, D., BRONSTEIN, A. M., BRONSTEIN, M. M., AND KIMMEL, R. 2010. Full and partial symmetries of non-rigid shapes. *Int. J. Comp. Vis.* 89, 18–39.
- RUSTAMOV, R. M. 2007. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proc. Symp. on Geom. Proc.*, 225–233.
- SHAMIR, A. 2008. A survey on mesh segmentation techniques. *Computer Graphics Forum* 27, 6, 1539–1556.
- SIMARI, P., KALOGERAKIS, E., AND SINGH, K. 2006. Folding meshes: hierarchical mesh segmentation based on planar symmetry. In *Proc. Symp. on Geom. Proc.*, 111–119.
- SUN, J., OVSJANIKOV, M., AND GUIBAS, L. J. 2009. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum (Special Issue of SGP)* 28, 5, 1383–1392.
- SURAZHISKY, V., SURAZHISKY, T., KIRSANOV, D., GORTLER, S. J., AND HOPPE, H. 2005. Fast exact and approximate geodesics on meshes. *ACM Trans. on Graph* 24, 3, 553–560.
- WANG, Y., XU, K., LI, J., ZHANG, H., SHAMIR, A., LIU, L., CHENG, Z., AND XIONG, Y. 2011. Symmetry hierarchy of man-made objects. *Computer Graphics Forum (Special Issue of Eurographics)* 30, 2, 287–296.
- XU, K., ZHANG, H., TAGLIASACCHI, A., LIU, L., LI, G., MENG, M., AND XIONG, Y. 2009. Partial intrinsic reflectional symmetry of 3D shapes. *ACM Trans. on Graph* 28, 5, 138:1–10.
- ZELNIK-MANOR, L., AND PERONA, P. 2004. Self-tuning spectral clustering. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, vol. 17, 1601–1608.