

Full and Partial Symmetries of Non-Rigid Shapes

Dan Raviv
Alexander M. Bronstein
Michael M. Bronstein
Ron Kimmel

Received: date / Accepted: date

Abstract Symmetry and self-similarity is the cornerstone of Nature, exhibiting itself through the shapes of natural creations and ubiquitous laws of physics. Since many natural objects are symmetric, the absence of symmetry can often be an indication of some anomaly or abnormal behavior. Therefore, detection of asymmetries is important in numerous practical applications, including crystallography, medical imaging, and face recognition, to mention a few. Conversely, the assumption of underlying shape symmetry can facilitate solutions to many problems in shape reconstruction and analysis. Traditionally, symmetries are described as extrinsic geometric properties of the shape. While being adequate for rigid shapes, such a description is inappropriate for non-rigid ones: extrinsic symmetry can be broken as a result of shape deformations, while its intrinsic symmetry is preserved. In this paper, we present a generalization of symmetries for non-rigid shapes and a numerical framework for their analysis, addressing the problems of full and partial exact and approximate symmetry detection and classification.

1 Introduction

“Symmetry, as wide or as narrow as you may define its meaning, is one idea by which man through the ages has tried to comprehend the created order, beauty, and perfection” [55]. These words of Hermann Weyl, one of the greatest twentieth century mathematicians, reflect the importance symmetry has in all aspects of our life. Symmetry, referred to in some contexts as *self-similarity* or *invariance* is the cornerstone of Nature, exhibiting itself through the shapes of

natural creations we see every day as well as through less evident yet omnipresent laws of physics.

The interest in symmetries of shapes dates back to the dawn of the human civilization. Early evidences that our predecessors attributed importance to symmetries can be found in many cultural heritages, ranging from monumental architecture of the Egyptian pyramids to traditional ancient Greek decorations. Johannes Kepler was among the first who attempted to give a geometric formulation to symmetries in his treatise *On the six-cornered snowflake* [25] in as early as 1611 (Figure 1). A few centuries later, the study of symmetric shapes became a cornerstone of crystallography. Finally, symmetries of more complicated higher-dimensional objects underlie modern physics theories about the nature of matter, space and time.

Since many natural objects are symmetric, symmetry breaking can often be an indication of some anomaly or abnormal behavior. Therefore, detection of asymmetries arises in numerous practical problems, among which medical applications are probably the first to come in mind. For example, detection of tumors in medical images can be based on deviations from otherwise symmetric body organs and tissues [35]. Facial symmetry is important in craniofacial plastic surgery [23], since symmetric facial features are often associated with beauty and aesthetics [38]. Furthermore, facial asymmetry can also be an indication of various syndromes and disorders [21]. Conversely, the assumption of symmetry can be used as a prior knowledge in many problems. It may facilitate, for example, the reconstruction of surfaces [53], face detection, recognition and feature extraction [44,49].

In pattern recognition and computer vision literature, there exists a significant number of papers dedicated to finding symmetries in images [36,51], two-dimensional [56,3,1] and three-dimensional shapes [54,24,41]. A wide spectrum of methods employed for this purpose includes approaches based on dual spaces [15], genetic algorithms [19], moments cal-

Department of Computer Science,
Tel.: +972-4-8294313
Fax: +972-4-8293900
Technion – Israel Institute of Technology, Haifa 32000, Israel
E-mail: {darav,bron,mbron,ron}@cs.technion.ac.il



Fig. 1 Hexagonal symmetry of ice crystals was one of the first shape symmetries to be rigorously studied (snowflake images taken by Wilson Bentley in 1902).

culuation [12], pair matching [33,14], and local shape descriptors [59].

Traditionally, symmetries are considered as *extrinsic* geometric properties of shapes that are related to the way the shape is represented in the Euclidean space. Though adequate for rigid shapes, such a point of view is inappropriate for non-rigid ones. Due to the deformations such shapes can undergo, the extrinsic symmetries may be lost, while *intrinsically* the shape still remains symmetric. Consider the human body example in Figure 2 (left). Extrinsic bilateral symmetry of the body is broken when the body assumes different postures (Figure 2, center). Yet, from the point of view of intrinsic geometry, the new shape remains almost identical, as such a deformation does not significantly change its metric structure. In this sense, intrinsic symmetries are a superset of the extrinsic ones.

An even more challenging problem is the detection of *partial* symmetries, shown in Figure 2 (right). In this example, the human figure has no leg, which makes only part of it symmetric. Detecting symmetric parts of generally asymmetric objects is a difficult problem in the rigid case [41], and significantly more difficult in the case of non-rigid shapes.

Recent works on deformable shape analysis studied many shape properties and characteristics remaining invariant under deformations in applications to shape similarity and correspondence. Anguelov *et al.* [2] addressed the problem of non-rigid shape correspondence based on local extrinsic properties. Fully intrinsic approaches were considered by Elad and Kimmel [16], Mémoli and Sapiro [40], and Bronstein *et al.* [9], who used the distortion of geodesic distances as a criterion of shape similarity. Reuter *et al.* [50] used the eigenvalues of the Laplace-Beltrami operator as shape descriptors, referred to as *Shape DNA*. Rustamov [52] modeled and compared shapes as distributions of commute time distances; a similar approach based on distributions of diffusion distances was presented by Mahmoudi and Sapiro [34]. Lévy [28] and Mateus *et al.* [37] used eigenmaps obtained by the first eigenfunctions of the Laplace-Beltrami operator

as low-dimensional Euclidean representations of non-rigid shapes.¹

Bronstein *et al.* [8,7,6] presented a framework for the computation of partial intrinsic similarity, where the similar parts are unknown in advance. The authors formulated a multi-criterion optimization problem in which the part “significance” and similarity are maximized at the same time; most similar and most significant parts are Pareto optima of the problem.

In [48], we introduced the notion of intrinsic symmetries for non-rigid shapes. Formulating non-rigid symmetries as intrinsic self-similarity allowed exploiting methods proposed for representation and comparison of non-rigid shapes. We used methods based on geodesic distances and motivated by [9], and presented several numerical tools for symmetry detection. In a parallel work, Ovsjanikov *et al.* [45] showed a spectral approach for intrinsic symmetry detection. The authors showed how reflection intrinsic symmetries are transformed into Euclidean ones in the space defined by the eigenfunctions of the Laplace-Beltrami operator. This approach is limited only to coping with reflection symmetries and cannot detect rotation and continuous symmetries. In addition, the method may be sensitive to geometric noise and, finally, it cannot be straightforwardly extended to dealing with partial symmetries. Yang *et al.* [58] showed an approach for the detection of reflection symmetries in 2D non-rigid shapes by finding axes maximizing the shape self-similarity.

Several generalizations of [48] for partial symmetries were proposed. Lasowski *et al.* [27] used a Markov random field model to obtain a probability distribution over all possible intrinsic matches of a shape to itself in order to reveal the symmetry structure. Xu *et al.* [57] used a voting procedure to find partial reflection symmetry axes and showed how the knowledge of symmetry can be exploited in shape segmentation and computer graphics applications.

¹ For additional methods, the reader is referred to [17,20].




Fig. 2 Symmetric or not? Visualization of the difference between extrinsic and intrinsic symmetry: an extrinsically symmetric shape is also intrinsically symmetric (left), however, an isometry of the shape is intrinsically symmetric but extrinsically asymmetric (center). The shape on the right, on the other hand, is partially intrinsically symmetric (the part obtained by removing the leg is symmetric).

1.1 Contributions

In this paper we elaborate on and expand the concepts put forward in [48]. Specifically, we classify and efficiently compute symmetries and partial symmetries, while using properties of symmetry groups in order to explore the symmetry space. Secondly, we generalize the notion of intrinsic symmetries to partial symmetries in the spirit of Bronstein *et al.* [6, 7] and show how partial symmetries can be found as a trade-off between self-similarity and partiality.

Compared to [45, 58] the main advantage of our approach for the detection of full symmetries are its ability to handle generic symmetries (not only reflections). In partial symmetry detection, compared to [57], our approach has a significantly lower computational complexity. Finally, using a generic metric framework, we have the possibility to use different metrics instead of the geodesic one.

The rest of this paper is organized as follows. In Section 2, we define intrinsic and extrinsic symmetries. In section 3, we introduce the space of approximate symmetries. Section 4 presents the relation between intrinsic and extrinsic symmetries and Section 5 deals with partial symmetries. Section 6 is devoted to a numerical framework for computation and visualization. Experimental results are presented in Section 7, and Section 8 concludes the paper.

2 Mathematical background

When dealing with nonrigid shapes, different geometric tools are invoked when the same shape is considered as a standalone rigid object or an instance (deformation) of a non-

rigid object. A unifying framework allowing to capture both points of view is possible by considering shapes from the perspective of *metric geometry* [16, 9, 7].

A geometric shape is modeled as a *metric space* (X, d) , where X is a two-dimensional smooth compact connected manifold (possibly with boundary) embedded into the Euclidean space \mathbb{E} (equal to \mathbb{R}^3 in case of three-dimensional objects and \mathbb{R}^2 in case of two-dimensional shapes), and $d : X \times X \rightarrow \mathbb{R}_+ \cup \{0\}$ is some *metric* measuring the distances on X . For the brevity of notation, we will write shortly X instead of (X, d) when the metric d is implied or not important.

There exist two most natural ways to define the metric d on X . One is to consider X as a subset of \mathbb{E} and measure the distances between points x, x' on X using the *restricted Euclidean metric*,

$$d_{\mathbb{E}}(x, x') = d_{\mathbb{E}|_{X \times X}}(x, x'). \quad (1)$$

The Euclidean metric regards the “external” properties of the shape, having to do with the way it is laid out in \mathbb{E} . We broadly refer to properties described by $d_{\mathbb{E}}$ as the *extrinsic geometry* of X .

Another way is to define the distance between x and x' as the length of the shortest path (*geodesic*) on the surface X connecting x and x' . We call the metric defined this way the *geodesic metric* and denote it by d_X . Properties defined by d_X are part of the *intrinsic geometry* of X . Broadly speaking, intrinsic geometry describes the properties of the shape which are invariant to inelastic deformations, that is transformations that do not stretch or tear the surface, while extrinsic geometry is associated with a specific rigid deformation. The same shape can be regarded both from the intrinsic

and extrinsic point of view by selecting d to be either the geodesic or the Euclidean metric, respectively [10, 39].

2.1 Intrinsic and extrinsic similarity

In order to say whether two shapes are similar, we compare them as metric spaces. From the point of view of metric geometry, two metric spaces are equivalent if their corresponding metric structures are equal. Such metric spaces are said to be *isometric*. More formally, given two metric spaces (X, d) and (Y, δ) , a bijective map $g : (X, d) \rightarrow (Y, \delta)$ is called an *isometry* if

$$\delta \circ (g \times g) = d. \quad (2)$$

In other words, an isometry is a metric-preserving map between two metric spaces, such that

$$d(x_1, x_2) = \delta(g(x_1), g(x_2)) \quad \forall x_1, x_2 \in X. \quad (3)$$

We call such (X, d) and (Y, δ) *isometric* and denote this by $(X, d) \sim (Y, \delta)$.

The definition of isometry obviously depends on the choice of the metric. Here, we consider two specific examples, the Euclidean metric $d_{\mathbb{E}}$ and the geodesic metric d_X . A bijection $g : (X, d_X) \rightarrow (Y, d_Y)$ satisfying $d_Y \circ (g \times g) = d_X$ is called an *intrinsic isometry*. Saying that (X, d_X) and (Y, d_Y) are isometric is equivalent to saying that X and Y are intrinsically similar.

On the other hand, if we consider the extrinsic geometry of the shapes (i.e., look at the shapes endowed with the Euclidean rather than geodesic metric), we notice that $(X, d_{\mathbb{E}})$ and $(Y, d_{\mathbb{E}})$ are subsets of the same metric space, $(E, d_{\mathbb{E}})$. As a result, an extrinsic isometry is a bijection between subsets of the Euclidean space rather than between two different metric spaces. In Euclidean geometry, the only possible isometries are rigid motions, which include rotation, translation and reflection transformations; we denote the family of such transformations by $\text{Iso}(\mathbb{E}, d_{\mathbb{E}})$. Thus, X and Y are extrinsically isometric if there exists $g \in \text{Iso}(\mathbb{E})$ such that $d_{\mathbb{E}}(X \times X) = d_{\mathbb{E}} \circ (g \times g)(X \times X)$. This means that two shapes are extrinsically isometric if one can be obtained by a rigid transformation of the other, which is sometimes expressed by saying that X and Y are *congruent*.

In the following, we will say that X and Y are *isometric* implying intrinsic isometry, and that X and Y are *congruent* when referring to an extrinsic isometry. The class of intrinsic isometries is usually richer than that of congruences, since any congruence is by definition also an intrinsic isometry. However, for some objects these two classes coincide, meaning that they have no incongruent isometries. Such shapes are called *rigid*, and their extrinsic geometry is completely defined by the intrinsic one². In particular, two-dimensional shapes realized as two-dimensional Euclidean

² More rigorously, the first fundamental form of a rigid shape defines (up to a congruence) its embedding into \mathbb{R}^3 .

sub-manifolds are always rigid³, unless they have point joints around which parts of the shapes can rotate [7, 30]. However, the assumption that the shape is a manifold rules out such singularities.

2.2 Symmetries

As mentioned in the Introduction, symmetries are self-similarities of shapes. So far, we have defined the rigorous meaning of similarity, using the notions of metric geometry. Self-similarity is a particular case, in which we compare a metric space to itself. A metric space (X, d) is self-similar if there exists a *self-isometry* on (X, d) (an isometry from (X, d) to itself).

However, just knowing that a shape is self-similar is not enough in order to understand how symmetric it is. For example, a sphere and a torus are self-similar, however, it is obvious that the sphere is “more symmetric” than the torus. We can actually claim that any shape is self-similar, since there always exists a trivial self-similarity, the identity transformation, which is an isometry by definition in any metric space.

A common and convenient way to model symmetries is by using the *group theory*, which describes operations between symmetries. A group, denoted by $(G, *)$ is a set with the associative binary operation $*$, satisfying:

- (G1) *Associativity*: $(g_1 * g_2) * g_3 = g_1 * (g_2 * g_3)$ for all $g_1, g_2, g_3 \in G$.
- (G2) *Identity*: There exists a unique identity element in G denoted by id , such that $g * \text{id} = g$ for all g in G .
- (G3) *Inverse*: For all g in G , there exists a unique inverse denoted by g^{-1} such that $g * g^{-1} = g^{-1} * g = \text{id}$.

A simple example is the group $(\mathbb{R}, +)$ of real numbers with the addition operation, in which the identity element is 0 and the inverse of a number is its negative. A subset of a group $(G' \subset G, *)$ with the operation $*$ restricted to G' and satisfying $g^{-1} * h \in G'$ for all $g, h \in G'$ is called a *subgroup* of $(G, *)$.

In our case, we are interested in groups as a concise description of classes of transformations acting on shapes. For this purpose, we consider the group $(\Pi(X), \circ)$ of all bijections $\Pi(X) = \{g : X \xrightarrow{1:1} X\}$ on the shape X with the function composition operator \circ . In this context, bijections from X to itself are referred to as *permutations* (even if X is continuous), as they can be thought of different ways to “permute” the points of the shape X . The identity element of the group

³ Two-dimensional shapes are manifolds restricted to the plane and therefore have a trivial second fundamental form. Isometries of a planar shape also have identical first fundamental forms, which by the fundamental theorem of the theory of shapes implies their congruence. This, in turn, implies rigidity of two-dimensional shapes.

is the identity map $\text{id}(x) = x$, and the inverse element is the inverse map g^{-1} .

Next, we recall that X is actually a metric space (X, d) equipped with the metric d and consider a subset of permutations $\Pi(X)$ which are also isometries,

$$\text{Iso}(X, d) = \{g \in \Pi(X) : d \circ (g \times g) = d\}. \quad (4)$$

It can be easily shown that $g^{-1}h \in \text{Iso}(X, d)$ for all $g, h \in \text{Iso}(X, d)$, which means that $(\text{Iso}(X, d), \circ)$ is a subgroup of $(\Pi(X), \circ)$. We refer to this subgroup as the *symmetry group* of (X, d) and use the notation

$$\text{Sym}(X, d) = (\text{Iso}(X, d), \circ), \quad (5)$$

omitting the operator \circ for notation brevity. Since this definition depends on the choice of the metric, we distinguish between the group of *extrinsic symmetries* $\text{Sym}(X, d_{\mathbb{E}})$ and that of *intrinsic symmetries* $\text{Sym}(X, d_X)$ as can be seen in Figure 3. Such a notation finally draws a formal distinction between these two notions.

It should be noted that in the case of rigid objects, these two notions coincide. Hence, in the 2D examples we show in the following, we refer to symmetries as both intrinsic and extrinsic ones.

The use of symmetry groups also allows to conveniently classify different types of symmetries. In many cases, the group describing the symmetries of a shape is isomorphic to a finite group, which can be thought of as a *representation* of the symmetry group. For example, the symmetries of an equilateral triangle (three rotations by 120 degrees and three reflections, see Figure 4) can be represented by the elements of the *dihedral group* D_3 , consisting of six elements R_0, R_1, R_2 (representing rotations) and S_0, S_1, S_2 (representing reflections), with the following composition rule (arranged as Cayley table),

\circ	R_0	R_1	R_2	S_0	S_1	S_2
R_0	R_0	R_1	R_2	S_0	S_1	S_2
R_1	R_1	R_2	R_0	S_1	S_2	S_0
R_2	R_2	R_0	R_1	S_2	S_0	S_1
S_0	S_0	S_2	S_1	R_0	R_2	R_1
S_1	S_1	S_0	S_2	R_1	R_0	R_2
S_2	S_2	S_1	S_0	R_2	R_1	R_0

which can be thought of as a ‘‘multiplication table’’ of the group.

Triskelion (Figure 5, right), a three-legged shape frequently occurring in ancient Greek ornaments and in modern heraldry, has three rotational symmetries described by the cyclic group C_3 , consisting of the cyclic permutations of the vector $(1, 2, 3)$. The Star of David (Figure 5, middle) is described by the dihedral group D_6 (six rotations and six reflections). A circle (Figure 5, left) has continuous symmetries, represented by the *special orthogonal group* $SO(2, \mathbb{R})$ (containing all the rotation transformations around a fixed point in a plane), and one reflective symmetry represented

by C_2 (the only group of order two). Combining them both creates $C_2 \times SO(2, \mathbb{R})$, which is the circle’s symmetry group. Note that the latter group is infinite – any infinitesimal rotation of the circle is a symmetry.

Another interesting symmetry groups worth mentioning are the *frieze groups* [32], defining one-parametric repetitive structures. There exist seven distinct frieze groups, generated by translation along one axis and a combination of rotations and reflections along another axis. In this paper, we do not distinguish between the different groups, yet one able to detect the repeating structure they produce.

The structure of the symmetry group (or a group isomorphic to it used as a representation) tells us how symmetric the shape is. If the group is trivial, that is, consists only of the identity element, the shape is said to be *asymmetric*. In general, when looking for symmetries we usually rule out the identity element. The symmetry analysis problem can therefore be split into two problems: finding the self-isometries of a shape (symmetry detection) and finding the group structure or the multiplication table (symmetry classification).

2.3 Symmetry generators

Given a group $(G, *)$ and a subset $G' \in G$, we denote by $\langle G' \rangle$ the *subgroup generated by G'* , defined as the smallest subgroup of G containing all the elements of G' . If $G = \langle G' \rangle$, i.e., every element of G can be expressed as the product of finitely many elements of G' and their inverses,

$$G = \{g_{i_1}^{\pm 1} * \dots * g_{i_K}^{\pm 1} : g_{i_1}, \dots, g_{i_K} \in G'\}, \quad (7)$$

we call G' the *generating set* of G and its elements the *generators* of G . If G' is finite, G is said to be *finitely generated*.

Applying this notion to symmetry group $\text{Sym}(X, d)$, we can find out that in many cases a few self-isometries can describe the entire symmetry structure of the shape. For example, the dihedral group D_3 visualized in Figure 4 has two generators: rotation by 120 degrees and reflection; other symmetries can be represented as finite compositions of these two. We call such generators of symmetry groups *symmetry generators*.

2.4 Approximate symmetries

The notion of the symmetry group allowed us to classify different shape symmetries and, in particular, answer the question whether one shape is more symmetric than another. Yet, perfect symmetry is a mathematical abstraction that never or rarely exists in natural shapes. Even the snowflakes we hailed as an example of symmetric shapes are never perfectly symmetric. In fact, according to our definition, the snowflakes from Figure 1 are asymmetric. Yet, with minor modifications, these shapes can be *symmetrized*.

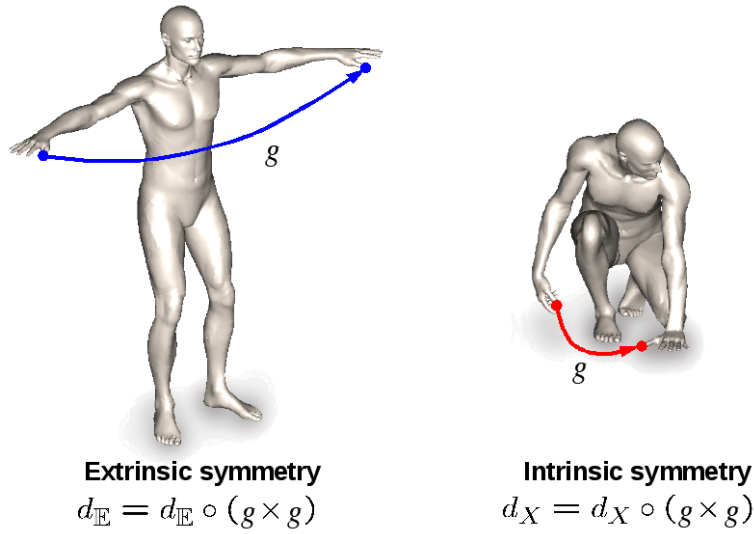


Fig. 3 Extrinsic versus intrinsic symmetry. Here, $g \in \Pi(X)$ is a permutation on X (a bijective map on X).

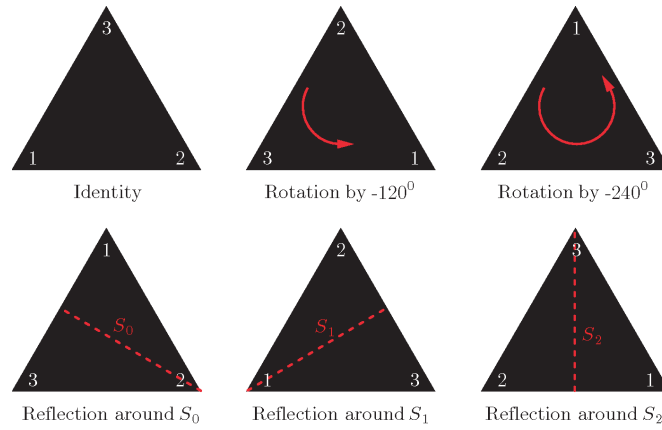


Fig. 4 Equilateral triangle has six symmetries (rotations and reflections), described by the dihedral group D_3 .




Fig. 5 Symmetries of different shapes can be described using group theory. Shown left-to-right: circle with continuous rotational symmetries and one reflective symmetry ($C_2 \times SO(2, \mathbb{R})$), star of David with dihedral symmetry (D_6 , including six rotations and six reflections) and triskelion with three rotational symmetries (described by the cyclic group C_3).

To this end, we introduce the notion of *approximate symmetry*: if we define a symmetry as the existence of a self-isometry on the metric space (X, d) , an approximate symmetry is the existence of an *approximate self-isometry* on

(X, d) . If a bijection $g : X \rightarrow X$ was said to be a self-isometry when $d \circ (g \times g) = d$ held, an approximate self-isometry should satisfy $d \circ (g \times g) \approx d$. Quantitatively, we measure how far g is from being an isometry by defining the *distor-*

tion

$$\begin{aligned} \text{dis}(g, d) &= \|d \circ (g \times g) - d\|_\infty \\ &= \sup_{x, x' \in X} |d(x, x') - d(g(x), g(x'))|. \end{aligned} \quad (8)$$

Since we assume compact spaces, the supremum is always achieved. In the following, we will be using maxima instead of suprema.

A map g with distortion $\text{dis}(g, d) \leq \varepsilon$ is called an ε -self-isometry. As a particular case, we have a self-isometry defined as a zero-self-isometry. We denote the family of all ε -self-isometries of (X, d) by

$$\text{Iso}_\varepsilon(X, d) = \{g \in \Pi(X) : \text{dis}(g, d) \leq \varepsilon\}. \quad (9)$$

As before, we can distinguish between extrinsic and intrinsic ε -self-isometries of (X, d) by defining the metric d to be $d_\mathbb{E}$ or d_X . Obviously, $\text{Iso}(X, d) \subset \text{Iso}_\varepsilon(X, d)$. The properties of approximate self-isometries are substantially different from those of exact self-isometries. If a composition of two self-isometries is still a self-isometry, composing two ε -self-isometries we get a 2ε -self-isometry. Consequently, $\text{Iso}_\varepsilon(X, d)$ is *not closed* under the function composition operation and thus *does not* form a group.

As a visualization of the consequences of the above differences, consider the torus shape depicted in Figure 6. The intrinsic symmetry group $\text{Sym}(X, d_X)$, consisting of all maps shifting the points along the tube (6, left), is isomorphic to $SO(2)$ and can be therefore parametrized by a single parameter (rotation angle). Since the torus has two reflective planes, and each reflection is isomorphic to C_2 , the intrinsic symmetry group is isomorphic to $C_2 \times C_2 \times SO(2)$. In this case the intrinsic symmetry group is isomorphic to the extrinsic one.

On the other hand, if some distortion is allowed, rotations across the tube can be considered as approximate symmetries. In this case, $\text{Iso}_\varepsilon(X, d_X)$ contains a two-parametric family of maps (6, right), in addition to the reflective ones.

2.5 Local and global asymmetry

We can calculate the shape's asymmetry both locally and globally. Usually, extrinsic asymmetry is calculated according to a reflective plane or a rotating vector. Another method to define asymmetry is based on the distortion of the symmetry as a function. Such a method is adequate for intrinsic symmetries as well.

In order to quantify how a point on X contributes to the asymmetry of the shape, we define the *local shape asymmetry*,

$$\text{asym}(X, x) = \max_{x' \in X} |d_X(x, x') - d_X(g^*(x), g^*(x'))| \quad (10)$$

quantifying the distortion of g^* at a point x . Points with large local asymmetry are responsible for symmetry breaking. The global asymmetry, with respect to g^* , can then be

written as,

$$\text{asym}(X) = \max_{x \in X} \text{asym}(X, x). \quad (11)$$

Using local asymmetry we can find local abnormality in intrinsically symmetric shapes.

3 Symmetry space

Though we cannot use group structures to represent approximate symmetry, we think of the space of permutation $\Pi(X)$, where each function has its distortion $\text{dis}(g, d)$. Approximate symmetries appear in this space as local minima of the distortion. The space of functions $\Pi(X)$ can also be endowed with a metric that measures the distance between two permutations of points on X . We define the metric between $f, g \in \Pi(X)$ as

$$d_{\Pi(X)}(f, g) = \max_{x \in X} d(f(x), g(x)) = d(f(X), g(X)), \quad (12)$$

which, in turn, depends on the choice of the metric d . We refer to the set

$$B_{\Pi(X)}(g, r) = \{f \in \Pi(X) : d_{\Pi(X)}(g, f) < r\} \quad (13)$$

as the *metric ball* (intrinsic or extrinsic, according to the choice of the metric in the definition of $d_{\Pi(X)}$) of radius r centered at g (we will omit r referring to a ball of some unspecified radius). A ball forms an open neighborhood of g . Since perfect symmetries have zero distortion, they are the global minimizers of the distortion on $\Pi(X)$. Moreover, they are also *local minimizers* of the distortion, in the sense that for every $g \in \text{Iso}_\varepsilon(X, d)$, there exists a sufficiently small neighborhood $B_{\Pi(X)}(g) \subset \Pi(X)$, such that any $f \in B_{\Pi(X)}(g)$ has $\text{dis}(f) \geq \text{dis}(g)$. We can therefore define approximate symmetries as

$$\begin{aligned} \text{Sym}_\varepsilon(X, d) &= \\ &\{g \in \text{Iso}_\varepsilon(X, d) : \text{dis}(g, d) \leq \text{dis}(f, d) \forall f \in B_{\Pi(X)}(g)\}. \end{aligned} \quad (14)$$

The exploration of the symmetry space consists of finding such local minima and composition relations between them. For increasingly large ε , we are likely to find more approximate symmetries, and for $\varepsilon = 0$ only perfect symmetries should be detected.

Though in the case of approximate symmetries there is no formal notion of generators (as there is no group structure), this idea can still be used for efficient exploration of the symmetry space. Given an initial set of possible approximate symmetries, we can search for a symmetry which is composed from two known symmetries. Each possible candidate is then compared to the known set to check if a new one was found, iterating until no more symmetries are added.

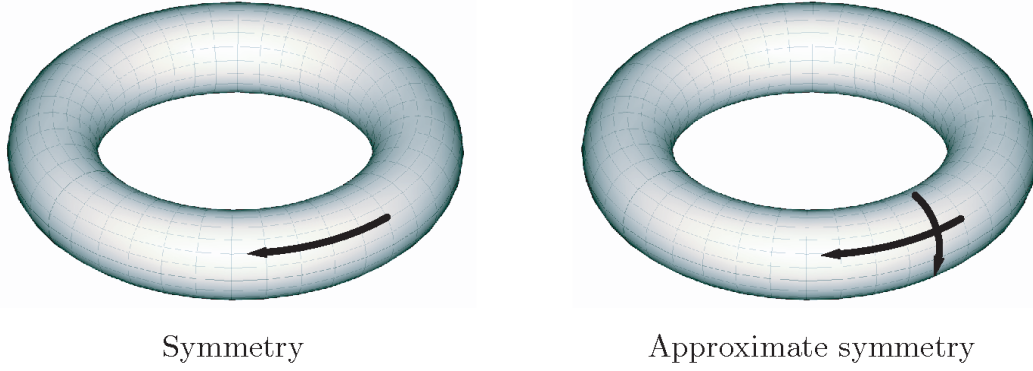


Fig. 6 Visualization of the difference between perfect and approximate symmetries. The intrinsic symmetry group of the torus consists of all maps rotating the points “along the tube” (left), and two reflections (vertical and horizontal). Approximate symmetries of the torus also include rotations “across the tube” (right).

4 Intrinsic symmetries as extrinsic symmetries

Let $(\mathbb{Z}, d_{\mathbb{Z}})$ be some homogenous metric space endowed with a simple metric (ideally, there should exist a closed form expression for $d_{\mathbb{Z}}$; we require homogeneity to obtain a simple isometry group $\text{Iso}(X)$). For a moment, let us also assume that there exists an *isometric embedding* $\varphi : (X, d_X) \rightarrow (\mathbb{Z}, d_{\mathbb{Z}})$ such that $d_X = d_{\mathbb{Z}} \circ (\varphi \times \varphi)$. We refer to the image $\varphi(X)$ as a *canonical form* of X in \mathbb{Z} [16]. Clearly, canonical forms are defined up to an isometry in \mathbb{Z} , since $d_{\mathbb{Z}} \circ (\varphi \times \varphi) = d_{\mathbb{Z}} \circ ((\varphi \circ i) \times (\varphi \circ i))$ for any $i \in \text{Iso}(\mathbb{Z})$. The canonical form $Z = \varphi(X)$ represents the intrinsic geometry of X in the sense that the two metric spaces (X, d_X) and $(Z, d_{\mathbb{Z}}|_Z)$ are isometric and, consequently, have isomorphic intrinsic symmetry groups.

Moreover, since the intrinsic geometry of Z coincides with its extrinsic counterpart, the analysis of the intrinsic symmetry group of the shape reduces to the analysis of the extrinsic symmetry group of its canonical form. Therefore, if the embedding space \mathbb{Z} has a reasonably simple isometry group (preferably with a convenient parametrization), the search for intrinsic symmetries is greatly simplified. For example, if $\mathbb{Z} = \mathbb{R}^3$, conventional extrinsic (Euclidean) symmetry detection algorithms can be employed [41].

This approach assumes the existence of an isometric embedding $\varphi : (X, d_X) \rightarrow (\mathbb{Z}, d_{\mathbb{Z}})$. Unfortunately, a perfectly isometric embedding does not exist in most cases. However, we can find the *minimum distortion embedding* of X into \mathbb{Z} ,

$$\begin{aligned} \varphi &= \arg \min_{\varphi: X \rightarrow \mathbb{Z}} \|d_X - d_{\mathbb{Z}} \circ (\varphi \times \varphi)\|_{\infty} \\ &= \arg \min_{\varphi: X \rightarrow \mathbb{Z}} \max_{x, x' \in X} |d_X(x, x') - d_{\mathbb{Z}}(\varphi(x), \varphi(x'))|, \end{aligned} \quad (15)$$

and repeat our reasoning replacing the assumption $d_X = d_{\mathbb{Z}} \circ (\varphi \times \varphi)$ with

$$\sup_{x, x' \in X} |d_X(x, x') - d_{\mathbb{Z}}(\varphi(x), \varphi(x'))| \leq \delta. \quad (16)$$

Proposition 1 *Let X be a shape, and let Z be its canonical form created by the embedding $\varphi : (X, d_X) \rightarrow (\mathbb{Z}, d_{\mathbb{Z}})$ with distortion δ . Then, for every $f \in \text{Iso}_{\varepsilon}(X, d_X)$, $\varphi \circ f \circ \varphi^{-1} \in \text{Iso}_{\varepsilon+2\delta}(\mathbb{Z}, d_{\mathbb{Z}})$; and for every $g \in \text{Iso}_{\varepsilon}(\mathbb{Z}, d_{\mathbb{Z}})$, $\varphi^{-1} \circ g \circ \varphi \in \text{Iso}_{\varepsilon+2\delta}(X, d_X)$.*

An alternative way to write Proposition 1 in terms of relations between the symmetry spaces is

$$\begin{aligned} \varphi \circ \text{Iso}_{\varepsilon}(X, d_X) \circ \varphi^{-1} &\subseteq \text{Iso}_{\varepsilon+2\delta}(\mathbb{Z}, d_{\mathbb{Z}}|_{Z \times Z}), \\ \varphi^{-1} \circ \text{Iso}_{\varepsilon}(\mathbb{Z}, d_{\mathbb{Z}}|_{Z \times Z}) \circ \varphi &\subseteq \text{Iso}_{\varepsilon+2\delta}(X, d_X). \end{aligned} \quad (17)$$

Observe that in the particular case of $\delta = 0$, the two spaces are equivalent; furthermore, if $\varepsilon = 0$, φ is a group isomorphism. We conclude that the applicability of intrinsic symmetry analysis based on canonical forms relies inherently on the ability to produce a low-distortion embedding φ . For example, if $\mathbb{Z} = \mathbb{R}^n$, the approach is suitable for nearly-flat shapes with small Gaussian curvature. If this is not the case then we cannot guarantee that intrinsic symmetries will be translated into extrinsic ones. In general, the canonical forms method is usually unsuitable for complicated intrinsic geometries, which cannot be faithfully represented as subsets of generic embedding spaces. For example, embedding a human body in different poses into \mathbb{R}^2 reveals the approximate reflective symmetry (Figure 7). Yet, embedding into \mathbb{R}^3 does not provide an extrinsic symmetric structure (Figure 8).

5 Partial symmetry

So far, our discussion assumed the existence of a bijection with zero or near-zero distortion, which we designated as a symmetry of the shape. In many cases, a shape does not have symmetries as a whole, yet possess parts that are symmetric. In order to extend our notion of intrinsic symmetries to this case, we need a definition of *partial similarity*.

Let $(X' \subseteq X, d|_{X' \times X'})$ be a *part* of the shape X , modeled as a metric sub-space of (X, d) with the restricted metric




Fig. 7 Embedding a human body into \mathbb{R}^2 reveals the reflective symmetry regardless of the body's pose.



Fig. 8 Embedding a human body into \mathbb{R}^3 fails to produce an extrinsically symmetric shape. The arms and legs are stretched in different directions.

$d|_{X' \times X'}(x, x') = d(x, x')$ for all $x, x' \in X'$. We denote by Σ_X the collection of all parts of X ,⁴ and by $p(X')$ the *partiality* of the part X' , a notion of the part significance with respect

to the entire shape, given e.g. by the relative area of $X \setminus X'$,

$$p(X') = 1 - \frac{\int_{X'} d\mu(x)}{\int_X d\mu(x)}. \quad (18)$$

Our definition of symmetry applied to the part X' is referred to as *partial symmetry*. Given a part X' with $p(X') = \lambda$, a symmetry (self-isometry) on $(X', d|_{X' \times X'})$ is a λ -*partial*

⁴ Formally, Σ_X is required to be a σ -algebra, see [6, 7].

symmetry. Since we do not know which part of the shape is symmetric, we need to look for the largest most symmetric part by minimizing the distortion under partiality constraints, i.e., look for λ -partial ε -self-isometric part with the smallest λ and ε . The two criteria are not commensurable, e.g., it is not clear what is better: a 0.5-partial 1-self-isometry or a 1-partial 0.5-self-isometry? Moreover, for any λ , there exists the identity map $\text{id} : X' \rightarrow X'$ with zero distortion.

We can thus define the set of partial approximate symmetries of X as

$$\text{Sym}_{\lambda,\varepsilon}(X, d) = \left\{ g \in \Pi(X') \left| \begin{array}{l} \text{dis}(g, d|_{X' \times X'}) \leq \varepsilon \\ \text{dis}(g, d|_{X' \times X'}) \leq \text{dis}(f, d|_{X' \times X'}) \\ \forall f \in \mathcal{B}_{\Pi(X')}(g) \\ p(X') \leq \lambda \end{array} \right. \right\} \quad (19)$$

Elements in $\text{Sym}_{\lambda,\varepsilon}(X, d)$ are called λ -partial ε -symmetries.

For a fixed partiality, the problem of finding the best partial symmetry is a scalar-valued constrained minimization problem,

$$\min_{\substack{X' \subset X \\ g \in \Pi(X')}} \text{dis}(g, d|_{X' \times X'}) \text{ s.t. } p(X') \leq \lambda_0, \quad (20)$$

In contrast to partial matching between two shapes, in order to rule out the trivial solution (identity map), we are not searching the global minimizer of (20) as done in [8]. Instead, we look for local minimizers of (20), which correspond to λ_0 -partial symmetries of X .

5.1 Regularization

As noted in [6], the straightforward definition of partiality (18) does not take into consideration the ‘‘quality’’ of the part and tends to produce multiple disconnected parts of X . As a remedy, in [6] it was proposed to add a *regularization* term, penalizing for the part boundary length,

$$r(X') = \int_{\partial X'} d\ell. \quad (21)$$

Using this definition given partiality $p(X') = p_0$ for a two-dimensional shape, the minimum is achieved by a circle. Unfortunately, no known extension exists for curved surfaces, since we may find two parts with similar area and boundary having an arbitrary number of disconnected components. [6] suggested another regularization based on Gauss-Bonnet theorem. Yet, we found the results based on parts length to be satisfactory. Adding a regularization term to our problem (20) yields

$$\min_{\substack{X' \subset X \\ g \in \Pi(X')}} \text{dis}(g, d|_{X' \times X'}) \text{ s.t. } \begin{cases} p(X') \leq \lambda_0; \\ r(X') \leq \rho_0. \end{cases} \quad (22)$$

Alternatively, one can move the regularization term to the objective function, obtaining

$$\min_{\substack{X' \subset X \\ g \in \Pi(X')}} \text{dis}(g, d|_{X' \times X'}) + \eta r(X') \text{ s.t. } p(X') \leq \lambda_0, \quad (23)$$

where η is the Lagrange multiplier governing the relative importance of the part regularity.

5.2 Fuzzy formulation

The main computational challenge in problems (20) and (22) is the need to perform optimization over all the subsets of X , which has combinatorial complexity. In [8], Bronstein *et al.* proposed a relaxation of the problem based on a *fuzzy approximation* of the parts. The part is represented as a *membership function* $u : X \rightarrow [0, 1]$, quantifying the probability of each point to belong to a part. The function u replaces X' in the above definitions, in the following way. The fuzzy distortion is defined as

$$\text{dis}(g, d) = \max_{x, x' \in X} u(x)u(x')|d(x, x') - (d \circ g)(x, x')|. \quad (24)$$

Note that u acts here as weight and the map $g \in \Pi(X)$ is a permutation on the entire X .

The fuzzy partiality is defined as

$$p(u) = \int_X (1 - u(x)) d\mu(x). \quad (25)$$

The regularization term, using a relaxation in the spirit of [43], is given by

$$r(u) = \int_X h(u(x)) \|\nabla_X u(x)\| d\mu(x), \quad (26)$$

where $h(t) \approx \delta(t - 0.5)$ is an approximation of the Dirac delta function, and $\nabla_X u$ is the intrinsic gradient of u .

The fuzzy version of (22) has the form

$$\min_{\substack{u: X \rightarrow [0, 1] \\ g \in \Pi(X)}} \max_{x, x' \in X} u(x)u(x')|d(x, x') - (d \circ g)(x, x')| \quad (27)$$

$$\text{s.t. } \begin{cases} p(u) \leq \lambda_0; \\ r(u) \leq \rho_0. \end{cases}$$

A fuzzy version of (23) is obtained in a similar way.

$$\min_{\substack{u: X \rightarrow [0, 1] \\ g \in \Pi(X)}} \max_{x, x' \in X} u(x)u(x')|d(x, x') - (d \circ g)(x, x')| + \eta r(u)$$

$$\text{s.t. } p(u) \leq \lambda_0. \quad (28)$$

6 Numerical framework

Even for shapes with simple intrinsic geometry, the complexity of its symmetry space is likely to be tremendous. The lack of a simple parametrization, similar to the one available for describing extrinsic symmetries, makes the analysis of symmetries of non-rigid shapes significantly more difficult.

Here, we propose an algorithm for automatic detection of symmetries comprising the following steps (see Figure 9): The input shape is first sub-sampled at a sparse set of sample points and an intrinsic descriptor is computed at each sample. Matches between similar descriptors are used to establish a set of initial correspondences, thus reducing the complexity of the search space. Next, a branch-and-bound procedure is used to prune correspondences with high distortion. The remaining coarse correspondences are refined using GMDS, and composition is performed to complete the group structure. The latter two stages are iterated until no new symmetries are found. In what follows, we describe each of the above steps in further detail.

6.1 Discretization and sampling

For simplicity, we assume the shape to be given in the form of a triangular mesh with N vertices; other discrete representations such as point clouds can be handled as well.

Intrinsic geometry is computed using the *fast marching* method [26], which produces a first-order approximation for the geodesic distances between points on the mesh.

Since the input sampling density is prohibitive for practical computation of symmetries, the mesh is sub-sampled. An R -sampling of the surface consisting of M points such set of points $X_R = \{x_1, \dots, x_M\} \subset X$ that form an R -covering, i.e., $X = \bigcup_{n=1}^M B_X(x_n, R)$, where B_X denotes a closed metric ball on X . A good sampling strategy of the surface can be achieved using the greedy *farthest point sampling* algorithm [16, 22, 42, 46], which guarantees that X_R is also R -separated, i.e., $d_X(x_i, x_j) \geq R$ for any $i \neq j$.

The coarse sampling X_R together with the $M \times M$ matrix of geodesic distances between each pair of samples form a discrete metric space, the set of permutations $\Pi(X_R)$ on which can be represented as M -tuples $g = (g_1, \dots, g_M) \in \{1, \dots, M\}^M$. Without loss of generality, we set $\pi_1 = (1, 2, \dots, M)$ to be the identity map.

6.2 Detection of coarse symmetries

Finding all π_k permutations with a distortion lower than ε requires computing the distortion of $\mathcal{O}(M!)$ mappings, which is prohibitive even for modest values of M . However, the search space can be reduced by ruling out mappings that are unlikely to have low distortion.

Following [18], we observe that in order for π to be a good candidate for an approximate symmetry, the intrinsic properties of the surface, such as the behavior of the metric d_X around every x_i should be similar to those around x_{π_i} . In order to quantify this behavior, for each $x_i \in X_R$ we compute the histogram $h_i = \text{hist}(\{\hat{d}_{ij} : \hat{d}_{ij} \leq \rho\})$ of the approximate

geodesic distances (\hat{d}_{ij}) in a ρ -ball centered at x_i . In our implementation, the parameter ρ was set to ∞ . The use of distance distributions is widely accepted in the literature. The reader is referred to the recent paper [31] for further discussion.

6.3 Local refinement

Once a coarse match is found it is used as an initialization for the second stage. We optimize over the images $x'_i = g(x_i)$ of a candidate symmetry g ,

$$\min_{x'_1, \dots, x'_N \in \hat{X}} \max_{i, j=1, \dots, N} |\hat{d}_{ij} - \hat{d}_X(x'_i, x'_j)|, \quad (29)$$

where the distance terms $\hat{d}_X(x'_i, x'_j)$ between arbitrary points on the mesh are found using the interpolation technique described in [9]. A local minimizer of (29) is found by convex optimization detailed in [5].

6.4 Partial symmetries

Solving (28) is done similarly to the framework presented in [6, 7]. We perform alternating minimization by first fixing u and solving for g and vice versa. u is initialized by the local asymmetry values of a candidate full symmetry and x' is initialized by interpolation.

For a fixed u , the minimization w.r.t. g is posed as a weighted GMDS problem,

$$\min_{x'_1, \dots, x'_N \in \hat{X}} \max_{i, j=1, \dots, N} u_i u_j |\hat{d}_{ij} - \hat{d}_X(x'_i, x'_j)|. \quad (30)$$

For a fixed g , we have the constrained problem

$$\begin{aligned} \min_{u_1, \dots, u_N} \max_{i, j=1, \dots, N} e_{ij} u_i u_j + \eta \sum_{i=1}^N h(u_i) a_i \sum_{k=1}^T q_{ik} g_k \\ \text{s.t. } \begin{cases} u_i \in [0, 1] & i = 1, \dots, N; \\ \sum_{i=1}^N (1 - u_i) a_i \leq \lambda_0, \end{cases} \end{aligned} \quad (31)$$

where $e_{ij} = |\hat{d}_{ij} - \hat{d}_X(x'_i, x'_j)|$ are fixed distortion terms, a_i are area elements,

$$q_{ik} = \begin{cases} \frac{1}{3} & \text{triangle } k \text{ shares the vertex } x_i, \\ 0 & \text{else,} \end{cases}, \quad (32)$$

$g_k = (\Delta u_k (X_k^T X_k)^{-1} \Delta u_k)^{1/2}$, with $X_k = (x_{k,2} - x_{k,1}, x_{k,3} - x_{k,1})$ being the 3×2 matrix with the local system of coordinates of the triangle k , and $\Delta u_k = (u_{k,2} - u_{k,1}, u_{k,3} - u_{k,1})$

Proposition 2 $g_k = |\nabla_X u_k|^2$.

See proof in Appendix A.

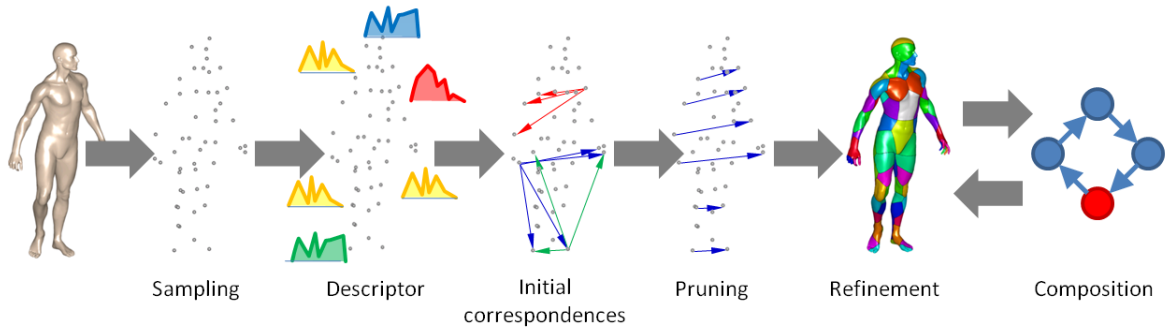


Fig. 9 Symmetry detection pipeline: The input mesh is sub-sampled at a sparse set of sample points and an intrinsic descriptor is computed at each sample. Matches between similar descriptors are used to establish a set of initial correspondences. Next, the branch-and-bound procedure is used to prune correspondences with high distortion. The remaining coarse correspondences are refined using GMDS, and composition is performed to complete the group structure. The latter two stages are iterated until no new symmetries are found.

6.5 Symmetry group completion by composition

The branch-and-bound procedure used for the detection of coarse symmetries is practical if the value of the threshold ϵ is relatively low. However, too low value of ϵ might result in rejecting true symmetries which due to acquisition and representation imperfections have high distortion. We observe that if the detected set of symmetries contains all the generators of the symmetry group, the missing group structure can be completed by their composition. If some of the generators are not detected, the completion will yield a subgroup of the symmetry group.

The completion algorithm proceeds as follows:

1. Input: set of refined symmetries $G = \{g_i\}$.
2. Compute all pair-wise compositions $h_{ij} = g_i \circ g_j$ for every $g_i, g_j \in G$.
3. For every h_{ij} having $\min_{g \in G} d_{\Pi(X)}(h_{ij}, g) > \delta$,
 - 3.1. Perform refinement of h_{ij}
 - 3.2. If $\text{dis } h_{ij} \leq \epsilon$, add h_{ij} to G .
4. Go to Step 2.

The procedure adds new low-distortion permutations resulting from a composition only if they lie at sufficient distance (controlled by the parameter δ) away from the already detected set of symmetries. It is guaranteed to stop after finite time, as in the worst case it will create a finite δ -separated covering of the bounded space $\Pi(X)$.

7 Results

In the following experiments, we show how to explore and visualize the space of intrinsic symmetries. We used triangular meshes from the TOSCA dataset [5], each consisting of 1000–2000 points. The branch-and-bound procedure was used in order to filter out strongly non-isometric permutations by measuring the score of matching between a relatively small number of local features (as features, we used

local histograms of geodesic distances). The branch-and-bound stage yielded a relatively small number of coarse-resolution permutations, which were considered as candidates for intrinsic symmetries. Refinement of these coarse mappings to establish high-resolution permutations was achieved using the GMDS procedure with the L_2 norm. For reasonable selections of ϵ , the execution time of the branch-and-bound step took a couple of seconds per surface on a 2.5GHz Intel CPU. The complexity of the GMDS-based refinement was about a minute.

7.1 Symmetry detection

Figure 10 (first row) presents the best two intrinsic symmetries of a human body. For this pose, the intrinsic and extrinsic symmetries are identical. Next, Figure 10 (second row) demonstrates how breaking the extrinsic symmetry of the body by changing its pose still preserves the intrinsic symmetry, as long as there is no considerable stretching of the limbs.

The same procedure can be applied for symmetry detection in planar shapes, which can be considered as a particular case of a flat surface with boundary. Figure 11 presents such a shape whose intrinsic symmetry is calculated w.r.t. the interior geodesic distances [30, 7].

Finally, Figure 12 presents a more complex set of ten symmetries of a five-legged octopus-like shape (a “pentapus”), and visualizes the symmetry composition approach for the exploration of the symmetry space of the shape. A perfectly symmetric “pentapus” would have a $C_2 \times C_5$ symmetry group (also known as D_5 or dihedral group of order five). Its generators are one rotation and one reflection. Since the deformation of the “pentapus” is not perfectly isometric, selecting too small an ϵ yields only a subset of D_5 . However, if the generators are in this subset, we can find the rest of the symmetries by composition. This is preferable over increasing the value of ϵ , which slows down the computation. Since




Fig. 10 Symmetries of the human shape. Top row: in this pose, the extrinsic and intrinsic symmetries are equivalent. The two self-isometries are identity and reflection (color represents corresponding points on the shape). Bottom row: in this pose, extrinsic symmetry is broken, yet the shape is still intrinsically symmetric. The two self-isometries are again identity and reflection.



Fig. 11 Symmetry of a planar shape w.r.t. to the interior geodesic metric. Like in the 3D case, the shape has two self-isometries: identity and reflection (color represents corresponding points on the shape).

the value of ε is unknown *a priori*, completion of the symmetry set by composition should always be performed as a part of the search procedure.

7.2 Comparison to Ovsjanikov *et al.*

In this section, we compare our approach to the algorithm of Ovsjanikov *et al.* [45]. This algorithm is based on the representation of the shape as an *eigenmap*

$$\Phi(x) = (\lambda_1^{-1/2} \phi_1(x), \dots, \lambda_K^{-1/2} \phi_K(x)), \quad (33)$$

where $\lambda_1, \dots, \lambda_K$ are non-repeating eigenvalues of the Laplace-Beltrami operator and ϕ_1, \dots, ϕ_K are the corresponding eigen-

functions. Ovsjanikov *et al.* [45] observe that any reflection symmetry $g \in \text{Sym}(X, d_X)$ satisfies $\phi_i \circ g = \pm \phi_i$ for $i = 1, \dots, K$. Thus, the symmetries of X can be parameterized by the *sign signature* $s = (s_1, \dots, s_K)$; $s_i \in \{-1, 1\}$ such that $\phi_i \circ g = s_i \phi_i$.

The symmetries of X are detected by testing different sign signatures. Given a sign signature s , define $\Phi_s(x) = (s_1 \lambda_1^{-1/2} \phi_1(x), \dots, s_K \lambda_K^{-1/2} \phi_K(x))$. Then,

$$E(s) = \sum_x \min_{x' \in X} \|\Phi_s(x) - \Phi_s(x')\|_2^2 \quad (34)$$




Fig. 12 Ten symmetries of the pentapus shape. A perfectly symmetric shape would have had a symmetry group $C_2 \times C_5$, generated by one rotation and one reflection. For each of the two possible reflections, the five rotations are presented in a separate row.

will vanish for s corresponding to intrinsic symmetries. For approximate symmetries, E is small.⁵ The symmetry itself is recovered as

$$g(x) = \operatorname{argmin}_{x' \in X} \|\Phi_s(x) - \Phi(x')\|_2. \quad (35)$$

The pentapus shape from Figure 12 was used for a comparison of our algorithm to Ovsjanikov *et al.* [45]. Eigenvalues and eigenvectors of the Laplace-Beltrami operator were computed as in [29] using the cotangent weights discretization [47]. We used sign signatures of length $K = 8$.

The smallest values of $E(s)$ were 0, obtained for the sequence of all pluses, and 15.7, obtained for the sequence $s = (-1, +1, +1, -1, +1, +1, -1, +1)$. The latter sequence encodes the reflection symmetry of the pentapus, while the former one encodes the identity map and all the rotation symmetries: they are all undistinguishable from each other (for a more detailed analysis, see the five-corner star example in Ovsjanikov *et al.* [45]). For other sign sequences $E(s)$ has significantly higher values of 46.88, 49.28, ... and they do not encode any of the desired symmetries of the shape. Thus, the algorithm of Ovsjanikov *et al.* [45] is able to find only two of the ten existing symmetries of the pentapus.

7.3 Exploring the symmetry space

Despite the high dimensionality of the space of permutations $\Pi(X)$, its metric structure can be visualized as a configuration of points in a low-dimensional Euclidean space, where each point represents a map in $\Pi(X)$, and the Euclidean distance between two points approximates $d_{\Pi(X)}$. Such an approximate representation can be constructed using multidimensional scaling [4]. An approximation of the distortion function is obtained by projecting the values of $\operatorname{dis}(g)$ onto its corresponding point in the representation space.

Figure 13 presents the approximate intrinsic symmetry set of a square with a bent corner. The square's extrinsic symmetry group is known to be generated by one rotation

⁵ This straightforward symmetry detection approach has complexity exponential in d . For large d , Ovsjanikov *et al.* [45] propose a fast heuristic.

and one reflection which creates the dihedral group D_4 . Bending one corner breaks most of the extrinsic symmetries (actually, only the identity and one reflection survive as extrinsic symmetries, which makes the extrinsic symmetry group of this shape isomorphic to C_2). Intrinsically, all eight symmetries survive the bending. Figure 13 visualizes these symmetries as clusters of low-distortion permutations in $\Pi(X)$.

Choosing the right ε obviously influences the solution we obtain. For $\varepsilon \approx 0$ only the identity mapping would be extracted, while choosing $\varepsilon \gg 1$, every permutation could be regarded as an approximate symmetry. In our experiments, choosing ε with a similar order as that of the resolution of the mesh produced good results for near-symmetric shapes. Figure 14 presents the influence of the value of ε on the set of symmetries. We varied ε from zero (left) to larger values (right) and obtained the identity member (left), the reflective symmetries (center) and finally a semi-uniform sampling of the whole space of permutations with bounded distortion (right). Approximate rotational symmetries can also be displayed as clusters. Figure 15 captures the ten clusters of the octopus symmetries embedded in \mathbb{R}^2 . Each cluster corresponds to a different symmetry.

The continuous symmetry set of a knot-shaped object is depicted in Figure 16. As the knot possesses a continuous family of rotation symmetries, the intrinsic symmetry group contains an infinite number of elements. Those are visualized as two two-dimensional contours of local minima (there are two of them due to an approximate reflective symmetry).

In Figure 17, the symmetries of the human body were embedded into \mathbb{R}^2 . The symmetry space was sampled in the vicinity of potential symmetries, and the distortion was interpolated over the entire domain. The color represents the value of asymmetry at each point. The four minima represent a reflection, half-reflection and their compositions.

Figure 18 presents the intrinsic reflection symmetry of a human body as a self-correspondence computed using GMDS. Given the computed symmetry $g(x_i) = x'_i$ and the groundtruth symmetry $g^*(x_i) = x_i^*$, we evaluated the accuracy of g comparing it to the groundtruth g^* ,

$$d_{\mathcal{G}}(g, g^*) = \frac{\sum_{i=1}^N d_X(x'_i, x_i^*)}{N \cdot \operatorname{Diam}(X)}, \quad (36)$$




Fig. 13 The eight elements of the square's D_4 dihedral group appear as clusters in the space of permutations approximately embedded into the Euclidean plane.




Fig. 14 The influence of ϵ on the approximate symmetry sets. From left to right: For $\epsilon = 0$ the identity member is extracted. For a small ϵ a set of approximate identity symmetries appear, and as ϵ is increased candidates for reflective symmetry start to emerge. Increasing ϵ even more, we obtain a semi-uniform sampling of the space of permutations with bounded distortions and the nice structures we experienced before disappears.

where $\text{Diam}(X)$ is the diameter of X , and N is the sample size. We achieved $d_{\mathcal{E}} \approx 0.009$ at three different resolutions, $N = 64, 128,$ and 256 . The computation took about 30, 70 and 260 seconds, respectively.

The refinement stage is performed on non-linear functions, which can converge to an undesired local minimum. We did not experience any difficulties in our experiments but

one can not guarantee a successful convergence for arbitrary initial conditions.

7.4 Local asymmetry

In the following experiment, we computed the *local shape asymmetry* of a human body with a local asymmetry that was introduced by elongating one of the arms. Figure 19




Fig. 15 The ten members of the $C_2 \times C_5$ symmetry group are shown as a clusters in the space of permutations approximately embedded into the Euclidean plane.




Fig. 16 The set of approximate intrinsic symmetries of a torus knot show up as two continuous contours in the space of permutations approximately embedded into \mathbb{R}^3 . Sliding the knot along itself shows up as a circle. One circle represent the reflection.

shows the local distortion of the detected reflection symmetry, which correctly localizes the deformed limb.

7.5 Partial symmetry

In order to demonstrate detection of partial symmetries, we used a female shape from the TOSCA dataset, whose approximate intrinsic symmetry was broken by removal of parts.

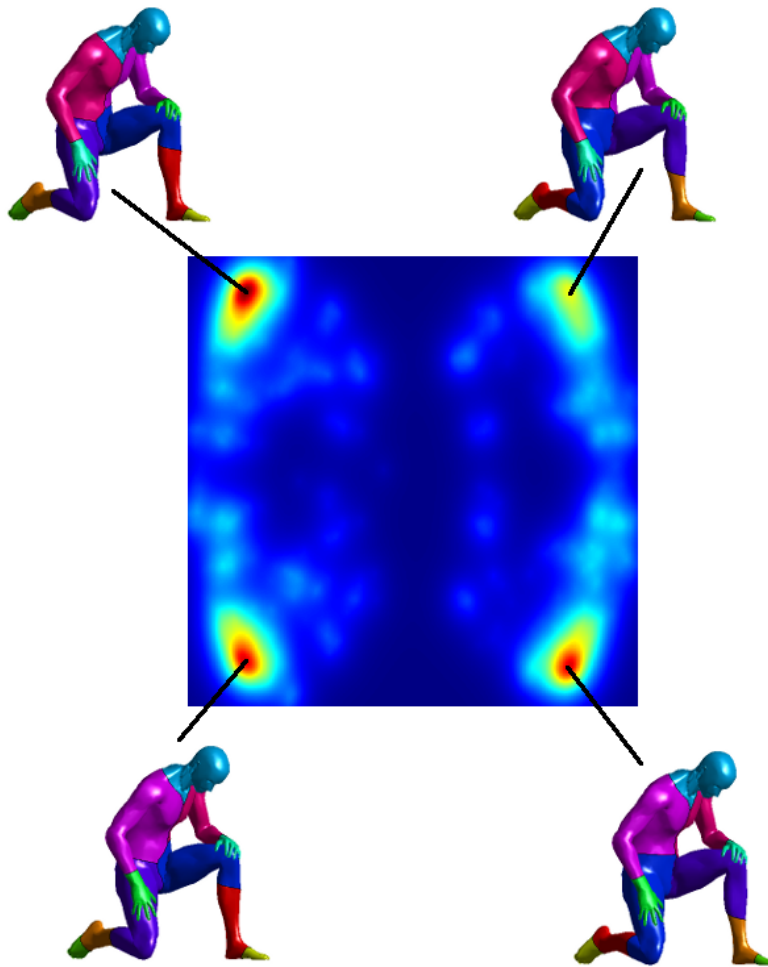


Fig. 17 The set of approximate intrinsic symmetries of a human body is embedded into \mathbb{R}^2 . Colors demonstrate normalized asymmetry values, where red represents symmetry and blue asymmetry. The four local minima correspond to the identity, full reflection, reflection of only half of the body, and their composition.




Fig. 18 Reflection symmetry of the human body depicted as a self-correspondence.



Fig. 19 Local distortion of the intrinsic reflection symmetry detected in a human shape whose intrinsic symmetry broken by deformation of one of arms. High distortion values are marked in red, correctly localizing the deformation.

We detected partial reflection symmetries by solving (28) with different values of partiality λ_0 and regularization coefficient η . The computation took about five minutes for each selection of the parameters. The obtained results are depicted in Figure 20. For visualization clarity, we completed the removed parts of the shape marking them in semi-transparent dark gray. Observe how the increase in the relative contribution of the regularization term (large η) tends to shorten the boundary of the detected part at the expense of its symmetry, while small values of η produce more symmetric and less regular parts. This phenomenon is particularly visible in the last two rows of Figure 20, where the detected part has multiple disconnected components disappearing with the increase of η . Figures 21 and 22 further visualize the shape of the selected part for different values of λ_0 and η . In all cases, the detected part appears to be more symmetric than the original shape. In Figure 23 and Figure 24 we depict more examples of partial symmetry detection. In Figure 23 we show the influence of partiality, and in Figure 24 the influence of regularization. As before, the detected parts appear more symmetric than the original shape. Since different coefficients of regularization and partiality provide different solutions, we can not predict, for a given shape, the best relation between them. More than that, an approximate full symmetry can be interpreted as a partial symmetry for different coefficients. In Figure 25 we present such a case. Once again, we can not determine a priori which solution is better.

8 Discussion and conclusions

We formulated the problem of approximate intrinsic symmetries detection which is specifically useful for non-rigid

articulated objects. The proposed measure of symmetry relies on the intrinsic geometric structure of the shape, namely the geodesic distances between surface points. It allowed us to find approximate intrinsic symmetries that are insensitive to bending of the shapes and detect and quantify asymmetric deformations. While other methods were presented recently, our method can also handle rotational symmetries and partial symmetries in sparse and dense sampling.

We presented a practical framework for the numerical computation of intrinsic symmetries, and demonstrated its potential by experimental results. We believe that the proposed approach could be useful for the detection of morphological distortions in medical imaging and we plan to explore its potential for diagnosis and analysis of morphometric deformations.

As a concluding remark, we emphasize that while the geodesic metric was used throughout this paper, the proposed framework is more general and is suitable for the detection of symmetries with respect to any metric. One of such possible alternatives is the *diffusion metric* [13], which is known to be significantly less sensitive to topological deformations than the geodesic counterpart. In [11], the GMDS framework was used to compute topologically-insensitive approximate isometries between shapes equipped with the diffusion geometry. In our future studies, we intend to develop this framework for the detection of full and partial symmetries.

Acknowledgements

This research was supported in part by the Israel Science Foundation (grant no. ISF 623/08), by The United States-

Israel Bi-national Science Foundation (grant no. BSF 2004274), and by the New York metropolitan research fund.

Appendix A

Proposition 2

$$g_k = |\nabla_X u_k|^2.$$

proof

For the proof of the proposition we will omit the triangle index k from X and u and use the subscript for partial derivation.

A linear discretization of the X and u leads to the local parameterization

$$X(w, v) = w \cdot (x_{k,2} - x_{k,1}) + v \cdot (x_{k,3} - x_{k,1}) \quad (37)$$

$$u(w, v) = w \cdot (u_{k,2} - u_{k,1}) + v \cdot (u_{k,3} - u_{k,1}),$$

for which the gradient in local coordinates becomes

$$u_w = u_{k,2} - u_{k,1} \quad (38)$$

$$u_v = u_{k,3} - u_{k,1}$$

$$X_w = x_{k,2} - x_{k,1}$$

$$X_v = x_{k,3} - x_{k,1}.$$

Hence, we can denote the local Riemannian metric as

$$\begin{bmatrix} E & F \\ F & G \end{bmatrix} = \begin{bmatrix} X_w \cdot X_w & X_w \cdot X_v \\ X_v \cdot X_w & X_v \cdot X_v \end{bmatrix} \quad (39)$$

Since the gradient of a function on a Riemannian manifold can be written in the local base as

$$\nabla_X u = \frac{u_w G - u_v F}{EG - F^2} X_w + \frac{u_v E - u_w F}{EG - F^2} X_v, \quad (40)$$

it follows that

$$\begin{aligned} |\nabla_X u|^2 &= \langle \nabla_X u, \nabla_X u \rangle = G \cdot u_w u_w - 2F \cdot u_w u_v + E \cdot u_v u_v \\ &= g_k. \end{aligned} \quad (41)$$

References

1. H. Alt, K. Mehlhorn, H. Wagnen, and E. Welzl. Congruence, similarity, and symmetries of geometric objects. *Discrete Comput. Geom.*, 3:237–256, 1988.
2. D. Anguelov, P. Srinivasan, H.C. Pang, D. Koller, S. Thrun, and J. Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Proc. NIPS*, 2005.
3. M. J. Atallah. On symmetry detection. *IEEE Trans. Computers*, c-34(7), July 1985.
4. I. Borg and P. Groenen. *Modern multidimensional scaling - theory and applications*. Springer-Verlag, New York, 1997.
5. A. Bronstein, M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer, 2008.
6. A. M. Bronstein and M. M. Bronstein. Not only size matters: Regularized partial matching of nonrigid shapes. In *Proc. Non-rigid Shape Analysis and Deformable Image Registration (NOR-DIA) workshop. See Proc. of Computer Vision and Pattern Recognition(CVPR)*, June 2008.
7. A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel. Analysis of two-dimensional non-rigid shapes. *International Journal of Computer Vision (IJCV)*, 78(1):67–88, June 2008.
8. A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel. Partial similarity of objects, or how to compare a centaur to a horse. *IJCV*, 84(2):163–183, 2009.
9. A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences (PNAS)*, 103/5:1168–1172, January 2006.
10. A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Rock, Paper, and Scissors: extrinsic vs. intrinsic similarity of non-rigid shapes. In *Proc. International Conference on Computer Vision (ICCV)*, 2007.
11. A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *IJCV*, 2009. to appear.
12. K.T. Cheung and H.S. Ip. Symmetry detection using complex moments. In *Proc. International Conference on Pattern Recognition (ICPR)*, volume 2, pages 1473–1475, 1998.
13. R.R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. Definition of diffusion distance.
14. H. Cornelius and G. Loy. Detecting rotational symmetry under affine projection. In *Proc. International Conference on Pattern Recognition (ICPR)*, volume 2, pages 292–295, 2006.
15. S. Derrode and F. Ghorbel. Shape analysis and symmetry detection in gray-level objects using the analytical fourier-mellin representation. *Signal Processing*, 84(1):25–39, 2004.
16. A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 25(10):1285–1295, 2003.
17. Ran Gal, Ariel Shamir, and Daniel Cohen-Or. Pose-oblivious shape signature. *IEEE Transactions on Visualization and Computer Graphics*, 13(2):261–271, 2007.
18. N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global registration. In *Proc. Symposium on Geometry Processing (SGP)*, pages 197–206, 2005.
19. Y. Gofman and N. Kiryati. Detecting symmetry in grey level images: The global optimization approach. In *Proc. International Conference on Pattern Recognition (ICPR)*, pages 951–956, 1996.
20. A.B. Hamza and H. Krim. Probabilistic shape descriptor for triangulated surfaces. In *Proc. IEEE International Conf. Image Processing (ICIP)*, volume 1, pages 1041–1044, September 2005.
21. S. Haraguchi, K. Takada, and Y. Yasuda. Facial Asymmetry in Subjects with Skeletal Class III Deformity. *Angle Orthodontist*, 72(1):28–35, 2001.
22. D. Hochbaum and D. Shmoys. A best possible heuristic for the k-center problem. *Mathematics of Operations Research*, 10:2:180–184, May 1985.
23. C.E. Huisinga-Fischer, J.P. Souren, F.S.B. Werken, B. Prahlandersen, and F. van Ginkel. Perception of Symmetry in the Face. *Journal of Craniofacial Surgery*, 15(1):128–134, 2004.
24. M. Kazhdan, B. Chazelle, D. Dobkin, T. Funkhouser, and S. Rusinkiewicz. A reflective symmetry descriptor for 3D models. *Algorithmica*, 38(1):201–225, 2003.
25. J. Kepler. *Strena seu de nive sexangula*. Frankfurt, 1611.
26. R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences (PNAS)*, 95(15):8431–8435, 1998.
27. R. Lasowski, A. Tevs, H.P. Seidel, and M. Wand. A Probabilistic Framework for Partial Intrinsic Symmetries in Geometric Data.
28. B. Levy. Laplace-beltrami eigenfunctions towards an algorithm that “understands” geometry. In *Proc. International Conference on Shape Modeling and Applications (ICSM)*, June 2006.

29. B. Lévy. Laplace-Beltrami eigenfunctions towards an algorithm that “understands” geometry. In *Int’l Conf. Shape Modeling and Applications*, 2006.
30. H. Ling and D. Jacobs. Shape classification using the inner-distance. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 29(2):286–299, 2007.
31. Rong F. Liu, Hao Zhang, Ariel Shamir, and Daniel Cohen-Or. A part-aware surface metric for shape analysis. volume 28, 2009.
32. Y. Liu, R. Collins, and Y. Tsin. A computational model for periodic pattern perception based on frieze and wallpaper groups. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 26(3):354–371, 2004.
33. G. Loy and J. Eklundh. Detecting symmetry and symmetric constellations of features. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 508–521, 2006.
34. M. Mahmoudi and G. Sapiro. Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models*, 71(1):22–31, 2009.
35. M. Mancas, B. Gosselin, and B. Macq. Fast and automatic tumoral area localisation using symmetry. In *Proc. International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, pages 725–728, 2005.
36. G. Marola. On the detection of axes of symmetry of symmetric and almost symmetric planner images. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 11(1), January 1989.
37. D. Mateus, R. Horaud, D. Knossow, and E. Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, June 2008.
38. L. Mealey, R. Bridgstock, and GC Townsend. Symmetry and perceived facial attractiveness: a monozygotic co-twin comparison. *Journal of Personality and Social Psychology*, 76(1):151–158, 1999.
39. F. Méholi. Gromov-hausdorff distances in Euclidean spaces. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, June 2008.
40. F. Méholi and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5:313–346, 2005.
41. N. J. Mitra, L. J. Guibas, and M. Pauly. Partial and approximate symmetry detection for 3D geometry. In *Proc. International Conference and Exhibition on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pages 560–568, 2006.
42. C. Moenning and N. Dodgson. A new point cloud simplification algorithm. In *Proc. International Conference on Visualization, Imaging and Image Processing*, 2003.
43. D. Mumford and J. Shah. Boundary detection by minimizing functionals. In *Proc. International Conference on Computer Vision (ICCV)*, 1990.
44. F. G. B. De Natale, D. D. Giusto, and F. Maccioni. A symmetry-based approach to facial features extraction. In *Proc. International Conference on Digital Signal Processing Proceedings (ICDSP)*, volume 2, pages 521–525, 1997.
45. M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. In *Proc. Eurographics Symposium on Geometry Processing (SGP)*, volume 27, 2008.
46. G. Peyré and L. Cohen. Geodesic remeshing using front propagation. *International Journal of Computer Vision (IJCV)*, 69(1):145–156, 2006.
47. U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, 2(1):15–36, 1993.
48. D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Symmetries of non-rigid shapes. In *Proc. Non-rigid Registration and Tracking (NRTL) workshop. See Proc. of International Conference on Computer Vision (ICCV)*, October 2007.
49. D. Reissfeld and Y. Yeshurun. Robust detection of facial features by generalized symmetry. In *Proc. International Conference on Pattern Recognition (ICPR)*, volume 1, pages 117–120, 1992.
50. M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace beltrami spectra as shape-dna of surfaces and solids. *Computer-Aided Design*, 38:342–366, 2006.
51. T. Riklin-Raviv, N. Sochen, and N. Kiryati. On symmetry, perspectivity, and level-set-based segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 31(8):1458–1471, 2009.
52. R. M. Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proc. Eurographics Symposium on Geometry Processing (SGP)*, pages 225–233, 2007.
53. I. Shimshoni, Y. Moses, and M. Lindernbaum. Shape reconstruction of 3D bilaterally symmetric surfaces. *International Journal of Computer Vision*, 39(2):97–110, 2000.
54. C. Sun and J. Sherrah. 3D symmetry detection using the extended gaussian image. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 19(2):164–168, 1997.
55. H. Weyl. *Symmetry*. Princeton University Press, 1983.
56. J.D. Wolter, T.C. Woo, and R.A. Volz. Optimal algorithms for symmetry detection in two and three dimensions. *The Visual Computer*, 1:37–48, 1985.
57. K. Xu, H. Zhang, A. Tagliasacchi, L. Liu, G. Li, M. Meng, and Y. Xiong. Partial intrinsic reflectional symmetry of 3d shapes. In *Proc. SIGGRAPH Asia*, 2009.
58. X. Yang, N. Adluru, L.J. Latecki, X. Bai, and Z. Pizlo. Symmetry of Shapes Via Self-similarity. In *Proc. International Symposium on Advances in Visual Computing*, pages 561–570, 2008.
59. H. Zabrodsky, S. Peleg, and D. Avnir. Symmetry as a continuous feature. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 17(12):1154–1166, 1995.

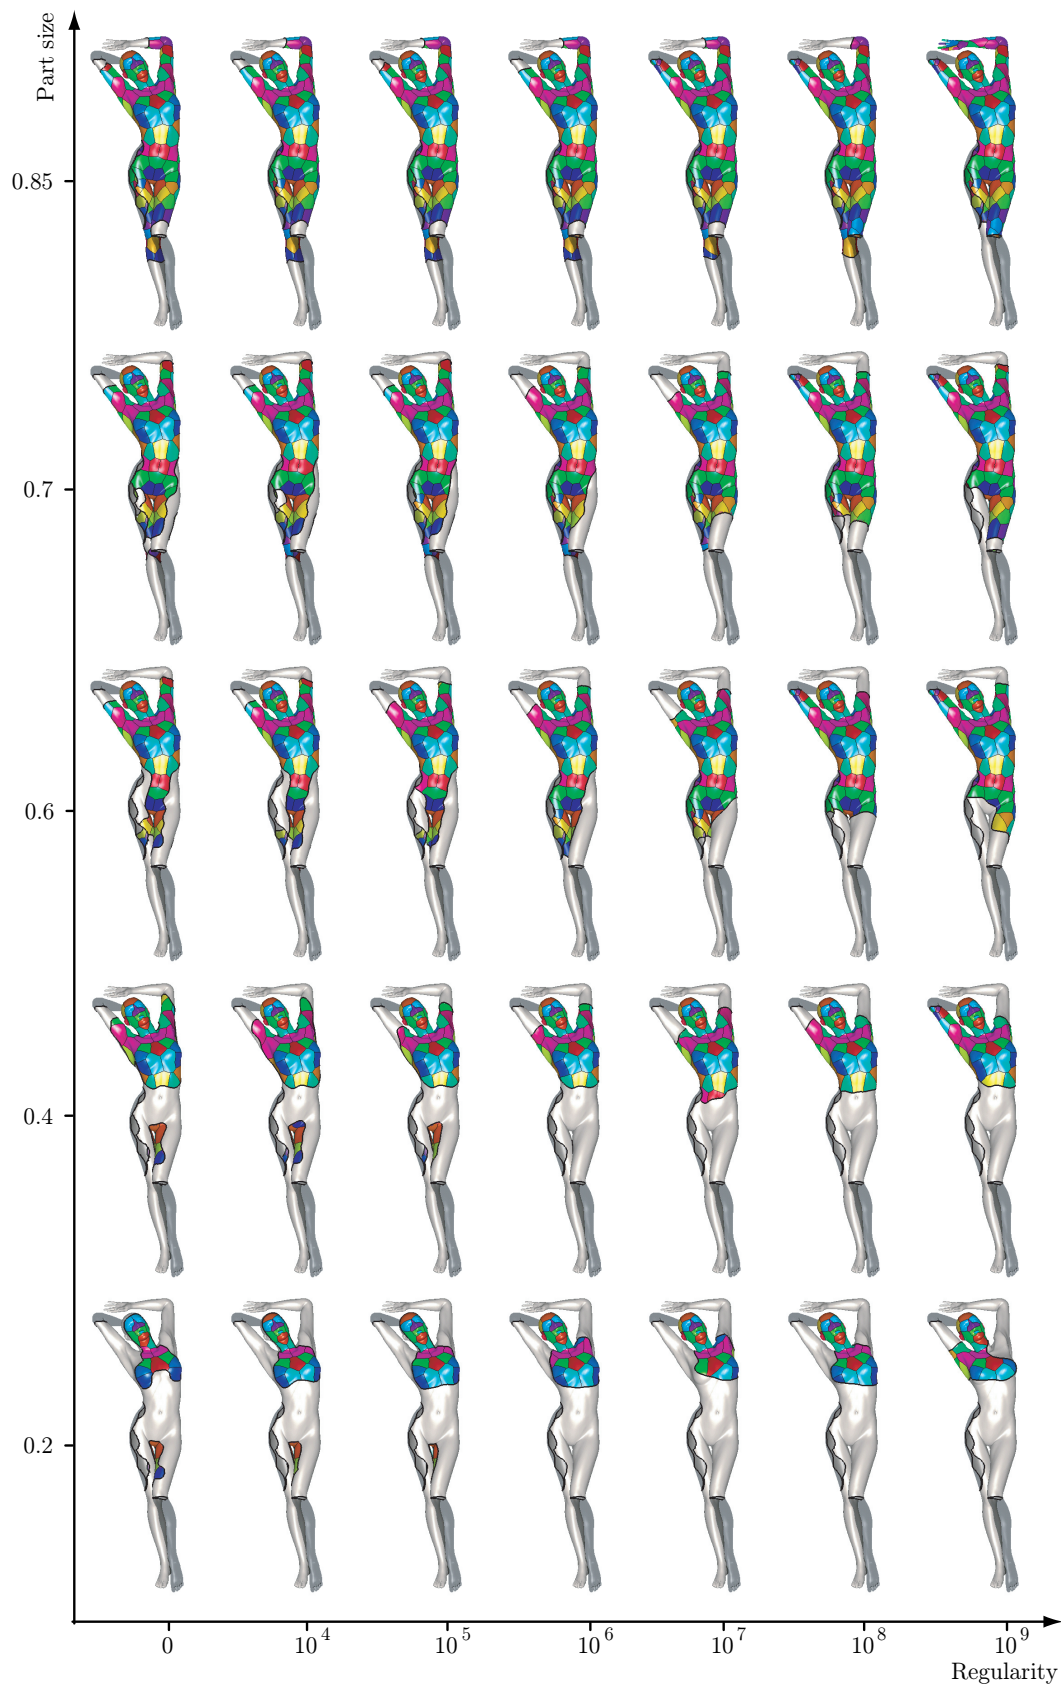


Fig. 20 Partial symmetry of a human body with broken intrinsic symmetry obtained by removal of parts (marked in semitransparent dark gray). The detected partial symmetries are shown as the function of the relative part size ($1 - \lambda_0$) and the regularization coefficient η . The discarded parts of the shape are marked in light gray. Colors encode corresponding regions. Note how the increase in η results in the shortening of the boundary at the expense of symmetry of the part (increase of $\text{dis}(X')$).




Fig. 21 Selected part for fixed $\eta = 10^4$ and part size varying from $\lambda_0 = 0.2$ (left) to 0.8 (right).




Fig. 22 Selected part for fixed $\lambda = 0.2$ and regularization coefficient varying from $\eta = 0$ (left) to 10^9 (right).




Fig. 23 Low partiality coefficient (left) versus a high one (right) in the pareto frontier. The symmetric surface is colored red.

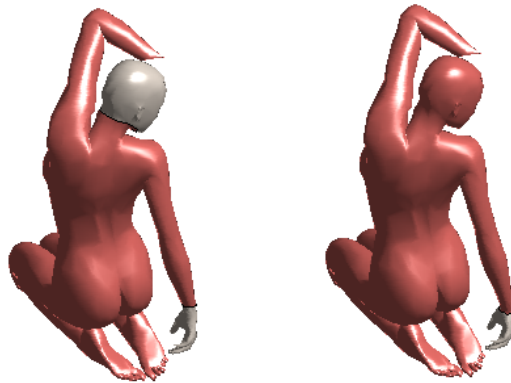


Fig. 24 Low regularization coefficient (left) versus a high one (right) in the pareto frontier. The symmetric surface is colored red.




Fig. 25 Ambiguity of partial symmetries: a shape with an asymmetric deformation can be interpreted in two ways: as a shape having an approximate full symmetry (left) or as a shape having an exact partial symmetry (right). Both interpretations correspond to Pareto-optimal choices of ϵ and λ .