Symmetries of non-rigid shapes

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

Department of Computer Science,
Technion – Israel Institute of Technology, Haifa 32000, Israel
{darav,bron,mbron,ron}@cs.technion.ac.il

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 pose the problem of finding intrinsic symmetries of non-rigid shapes and propose an efficient method for their computation.

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" [29]. These words of Hermann Weyl, one of the greatest twentieth century mathematician, 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* [15] in as early as 1611. 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.

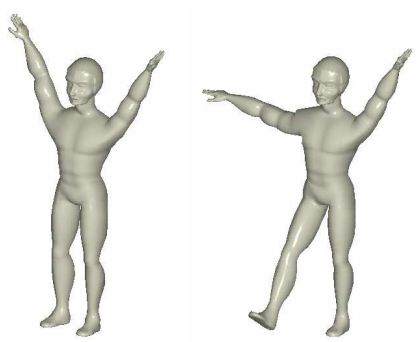

Figure 1. Symmetric or not? Visualization of the difference between extrinsic and intrinsic symmetry: extrinsically symmetric shape is also intrinsically symmetric (left), however, an isometry of the shape is intrinsically symmetric but extrinsically asymmetric (right).

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 [18]. Facial symmetry is important in craniofacial plastic surgery [13], since symmetric facial features are often associated with beauty and aesthetics [20]. Further-

*This research was partially supported by the Israeli Science Foundation (ISF) grant No. 738/04, the United States Israel Binational Science Foundation (BSF) grant No. 2002474 and by the Horowitz Fund.

more, facial asymmetry can also be an indication of various syndromes and disorders [12]. Conversely, the assumption of symmetry can be used as a prior knowledge in many problems. It may facilitate, for example, the reconstruction of surfaces [27], face detection, recognition and feature extraction [24, 26].

In pattern recognition and computer vision literature, there exists a significant number of papers dedicated to finding symmetries in images [19], two-dimensional [30, 2, 1] and three-dimensional shapes [28, 14, 22]. A wide spectrum of methods employed for this purpose includes approaches based on dual spaces [7], genetic algorithms [11], moments calculation [5], pair matching [17, 6], and local shape descriptors [31].

Traditionally, symmetries are considered as *extrinsic* geometric properties of shapes, i.e., 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 or deformable ones. Due to the deformations such shapes can undergo, the extrinsic symmetries may be lost, while *intrinsically* the shape still remains symmetric. Considering as an example the human body in Figure 1 (left). Extrinsic bilateral symmetry of the body is broken when the body assumes different postures. 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. Considering intrinsic rather than extrinsic symmetries allows us to characterize the object self-similarity that is invariant to deformations. If we resort to our previous medical illustration, we can detect tumors as irregular, non-symmetric objects, no matter how the tissues and the symmetric organs surrounding them are bent.

Unfortunately, while extrinsic symmetry computation is a well-established subject, very little effort has so far been devoted to intrinsic symmetries. Similarity of non-rigid shapes was considered in recent papers of Elad and Kimmel [8], Mémoli and Sapiro [21], and Bronstein *et al.* [3]. Yet, none of these papers deals with self-similarity and symmetry. In this paper, we address the problem of finding intrinsic symmetries of non-rigid shapes. We define criteria of global and local asymmetry and present a practical method for their computation. The rest of this paper is organized as follows. In Section 2 we define intrinsic and extrinsic symmetries. Section 3 is devoted to a numerical scheme for symmetry computation.Experimental results are presented in Section 4. Section 5 concludes the paper.

2. Extrinsic and intrinsic symmetries

We model a non-rigid shape X as a two-dimensional smooth compact connected and complete Riemannian surface (possibly with boundary) embedded into $\mathbb{R}^3$. The manifold can be considered as a metric space with the *geodesic metric* $d_X : X \times X \to \mathbb{R}$ measuring the lengths of the shortest paths on X, where the length structure is induced by the Euclidean metric $d_{\mathbb{R}^3}$. Another metric, with which X can be equipped, is the *restricted Euclidean metric* $d_{\mathbb{R}^3}|_X$. We broadly refer to properties described in terms of the metric d_X as to the *intrinsic geometry* of X, and to properties associated with $d_{\mathbb{R}^3}|_X$ as the *extrinsic geometry*.

A transformation $g : X \to \mathbb{R}^3$ preserving the extrinsic geometry of a shape X is called a *congruence*, and X and $g(X)$ are called *congruent*. If in addition $g(X) = X$, such a g is called a *self-congruence* or an *extrinsic symmetry* of X. Extrinsic symmetries form a group with the function composition,

$$\mathrm{ESym}(X) = \{g : X \overset{1:1}{\to} X : d_{\mathbb{R}^3}|_X = d_{\mathbb{R}^3}|_X \circ (g \times g)\},$$

which we call the *extrinsic symmetry group* of X. $\mathrm{ESym}(X)$ contains a subset of rotation and reflection transformations in $\mathbb{R}^3$, to which the shape X is invariant.

Analogously, a transformation $g : X \to \mathbb{R}^3$ preserving the intrinsic geometry of a shape X is called an *isometry*, and the shapes X and $g(X)$ are referred to as *isometric*. An isometry mapping X onto itself is called a *self-isometry* or an *intrinsic symmetry* of X. Like their extrinsic counterparts, intrinsic symmetries form the *intrinsic symmetry* or the *isometry group*

$$\mathrm{ISym}(X) = \{g : X \to X : d_X = d_X \circ (g \times g)\}. \quad (1)$$

The intrinsic symmetry group is invariant to isometries of the shape. Since every extrinsic symmetry is also an intrinsic symmetry, $\mathrm{ESym}(X)$ is a subset of $\mathrm{ISym}(X)$. This fact is visualized in Figure 1, showing that an extrinsically symmetric shape is also intrinsically symmetric, yet an intrinsically symmetric shape might be extrinsically asymmetric. In the sequel, we will focus our attention on intrinsic symmetries only, to which we will refer simply as to "symmetries".

We denote by $\mathbb{F}(X)$ the space of all mappings $g : X \to X$ (not necessarily bijective), and define a metric on $\mathbb{F}(X)$ as

$$d_{\mathbb{F}(X)}(f, g) = \max_{x \in X} d_X(f(x), g(x)) \quad (2)$$

for all $f, g \in \mathbb{F}(X)$ (note that the maximum is always achieved, since X is a compact space). We refer to the set

$$B_{\mathbb{F}(X)}(g, r) = \{f \in \mathbb{F}(X) : d_{\mathbb{F}(X)}(g, f) \leq r\} \quad (3)$$

as to the *closed metric ball* of radius r centered at g (as a matter of notation, we will omit r referring to a ball of some unspecified radius). A ball forms a closed neighborhood of g. We also associate the *distortion*

$$\mathrm{dis}(g) = \max_{x, x' \in X} |d_X(x, x') - d_X(g(x), g(x'))| \quad (4)$$

with every mapping $g \in \mathbb{F}(X)$. In this context, the distortion is a function $\mathrm{dis} : \mathbb{F} \to \mathbb{R}$ measuring how d_X differs from $d_X \circ (g \times g)$.

In these terms, an intrinsic symmetry can be defined as a mapping $g \in \mathbb{F}(X)$ having $\mathrm{dis}(g) = 0$, which is a *global minimum* of dis. Moreover, it can be shown that intrinsic symmetries are also *local minima* of the distortion, in the sense that for every $g \in \mathrm{ISym}(X)$, there exists a sufficiently small neighborhood $B_{\mathbb{F}(X)}(g) \subset \mathbb{F}(X)$, such that any $f \in B_{\mathbb{F}(X)}(g)$ has $\mathrm{dis}(f) \geq \mathrm{dis}(g)$.

2.1. Approximate symmetries

In real applications, due to acquisition and representation inaccuracies, perfect symmetry rarely exists. Non-symmetric shapes have a trivial intrinsic symmetry group, containing only the identity mapping $\mathrm{id}(x) = x$. However, while not symmetric in the strict sense, a shape can still be *approximately symmetric*. An intuitive way to understand the difference between the two definitions, is by thinking of a non-symmetric shape as obtained by applying a deformation to some other symmetric shape. Such a deformation may break the symmetries of the shape: if previously a symmetry was a self-isometry, we now have mappings which have non-zero distortion.

We call a map $g \in \mathbb{F}(X)$ with $\mathrm{dis}(g) \leq \epsilon$ an *ϵ-self-isometry*, and, by analogy with symmetries, define the *intrinsic ϵ-symmetries* of X as the set of ϵ-self-isometries which are local minima of dis on $\mathbb{F}(X)$,

$$\mathrm{ISym}(X, \epsilon) = \left\{ g \in \mathbb{F}(X) : \epsilon \geq \mathrm{dis}(g) = \min_{f \in B_{\mathbb{F}(X)}(g)} \mathrm{dis}(f) \right\}. \tag{5}$$

Note that unlike symmetries, ϵ-symmetries are not closed under function composition and therefore do not form a group.

Also note that in practice, ϵ is unknown *a priori*, and choosing different values of ϵ results in different structure of $\mathrm{ISym}(X, \epsilon)$. Experiments show that choosing ϵ to be twice the sampling radius of discrete shapes produces meaningful results.

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

$$\mathrm{asym}(X, x) = \max_{x' \in X} |d_X(x, x') - d_X(g^*(x), g^*(x'))| \tag{6}$$

measuring the distortion of g^* at a point x. Points with large local asymmetry are responsible for symmetry breaking.

The global assymetry,with relation to g^*, can be written as,

$$\mathrm{asym}(X) = \max_{x \in X} \mathrm{asym}(X, x). \tag{7}$$

In practice, the use of an L_p version of the distortion in lieu of its L_∞ counterpart appears to be more robust to noise.

3. Computation of intrinsic symmetries

For practical computation of symmetries, the surface X has to be discretized and sampled at N points, constituting an *r-covering* (i.e., $X = \bigcup_{n=1}^{N} B_X(x_n, r)$, where B_X denotes a closed metric ball on X). We denote this sampling by $X_r = \{x_1, ..., x_N\} \subseteq X$. A good sampling of the surface can be achieved using the *farthest point sampling* algorithm [8, 9, 23, 25], 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 extrinsic geometry of X is approximated by a triangular mesh $\hat{X}$ built upon the vertices X_r. In order to approximate the intrinsic geometry, we use the fast marching method [16], which produces a first-order approximation for the geodesic distances between X_r on $\hat{X}$.

Our goal is to approximate (7) and (6) on the discrete shape, or in other words, find such a mapping that has minimum distortion among all the non-trivial local minima of dis. In general, this is an NP-hard problem involving combinatorial optimization [21]. In [3], Bronstein *et al.* proposed a relaxation for problems of this type by means of continuous optimization on the triangular mesh $\hat{X}$, in the spirit of the multidimensional scaling (MDS) problem. Their framework is referred to as *generalized MDS* (GMDS).

Here, we adopt this approach and actually approximate (7) and (6) on $\mathbb{F}(\hat{X})$ rather than $\mathbb{F}(X_r)$. The method proposed here consists of two stages: first, an initialization is found by computing the non-trivial minimum-distortion discrete mapping on a sub-sampling of the shape. Next, the initialization is refined by means of GMDS, which finds a local minimizer of the distortion.

3.1. Coarse initialization

Given X_r, we sub-sample it with a larger radius R, producing a sparser sampling $X_R \subset X_r$ containing $M \ll N$ points. We denote by $\mathbb{F}(X_R)$ the set of all mappings $\pi : X_R \to X_R$ (permutations on the discrete set X_R), which can be represented as M-tuples $\pi = (\pi_1, ..., \pi_M) \in \{1, ..., M\}^M$. Without loss of generality, we set $\pi_1 = (1, 2, ..., M)$ to be the identity map.

Finding the global minimum on $\mathbb{F}(X_R) \setminus \{\pi_1\}$ potentially requires computing the distortion of $\mathcal{O}(M^M)$ mappings. However, the search space can be reduced by ruling out mappings that are unlikely to have low distortion. 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 $h_i = \mathrm{hist}(\{\hat{d}_{ij} : \hat{d}_{ij} \leq \rho\})$ of the approximate geodesic

distances in a ρ-ball centered x_i (in our implementation, the parameter ρ was set to ∞).

Using the vectors h_i as local descriptors of the points in X_R, we compute the dissimilarity of two points $x_i, x_j \in X_R$. Our practice shows that a straightforward use of the Euclidean distance $\|h_i - h_j\|_2$ between the descriptors may be inaccurate due to binning errors. To account for distances between the bins, we use the weighted histogram distance, $d(h_i, h_j) = \sqrt{(h_i - h_j)^{\mathrm{T}} A (h_i - h_j)}$, where A_{mn} are the distances between the bins m and n. For each point x_i in X_R, we construct a set $C_i \subset \{1, ..., M\}$ of indices of K points having the most similar descriptors. K is selected to be a small number, typically significantly smaller than M. We define the reduced search space $\mathbb{F}_{\mathrm{init}} = C_1 \times C_2 \times ... \times C_M$. Mappings copying any x_i to $x_{\pi_i} \notin C_i$ are excluded from the search space. Since $i \notin C_i$, the identity mapping is also excluded from $\mathbb{F}_{\mathrm{init}}$.

Even though the coarse sample size M and the number of initial matches for every point are relatively small, $\mathbb{F}_{\mathrm{init}}$ has still $\mathcal{O}(K^M)$ mappings, making an exhaustive search prohibitively expensive. However, we can select a reasonably good mapping from $\mathbb{F}_{\mathrm{init}}$ using the following hierarchical greedy algorithm, which is similar in its spirit to the technique presented in [10] for improving convergence of ICP-based extrinsic surface alignment.

1. **Pairing:** For each pair $(i, j) \in \{1, ..., M\}^2$, choose the best pair $(m, n) \in C_i \times C_j$ minimizing the distortion $\left|\hat{d}_{ij} - \hat{d}_{mn}\right|$. This establishes a two-point correspondence $(i, j) \mapsto (m, n)$. The outcome of this step is the set of $\mathcal{O}(M^2)$ two-point correspondences E_2, which we sort in increasing order of distortion.

2. **Merging:** The pairs are merged into four-point correspondences. Taking the first two-point correspondence $e \in E_2$, we find another two-point correspondence having a disjoint domain and minimizing the distortion of the obtained four-point correspondence. We remove all correspondences sharing the same domain from E_2 and continue until E_2 becomes empty. The merging continues hierarchically, producing E_{2k} from E_k, stopping typically at E_8 or E_{16}.

3. **Completion:** We select the minimum distortion correspondence $(i_1, ..., i_k) \mapsto (\pi_{i_1}, ..., \pi_{i_k})$ from the last produced E_k, and complete it to a full M-point correspondence by adding the missing indices $\{i_{k+1}, ..., i_M\} = \{1, ..., M\} \setminus \{i_1, ..., i_k\}$ and their images $\pi_{i_{k+1}}, ..., \pi_{i_M}$. For each added point j, we select

$$\pi_j \quad = \quad \arg \min_{\pi_j \in \{1,...,M\}} \max_{i \in \{i_1,...,i_k\}} \left|\hat{d}_{ij} - \hat{d}_{\pi_i, \pi_j}\right|.$$

Return the found mapping $\pi_{\min}$ and its distortion $\epsilon_{\min}$.

Since the algorithm never backtracks, it may produce a suboptimal mapping $\pi_{\min}$. However, practice shows that if some good pairs are found at Step 1, the algorithm tends to produce a very good estimate for the minimum distortion mapping on $\mathbb{F}_{\mathrm{init}}$.

A guaranteed global minimum on $\mathbb{F}_{\mathrm{init}}$ can be computed by using a branch and bound algorithm, similar to [10]. The idea of the algorithm is based on the fact that if a good estimate for $\pi_{\min}$ is found using the greedy matching, a large set of mappings in $\mathbb{F}_{\mathrm{init}}$ can be eliminated efficiently. The algorithm is initialized by $\pi_{\min}$ and $\epsilon_{\min}$ found by greedy matching, and proceeds as following.

1. Given a correspondence of $k-1$ feature points $(1, ..., k-1) \mapsto (\pi_1, ..., \pi_{k-1})$, we would like to establish $k \mapsto \pi_k$.

2. **Prune:** For each potential correspondence $\pi_k \in C_k$, evaluate $\max_{i=1,...,k} |\hat{d}_{ik} - \hat{d}_{\pi_i, \pi_k}|$. If the obtained distortion is larger than $\epsilon_{\min}$, discard the potential correspondence.

3. **Branch:** For each remaining π_k, recursively invoke Step 1 with $(1, ..., k) \mapsto (\pi_1, ..., \pi_k)$.

4. **Bound:** If $k = M$, compute the distortion $\mathrm{dis}\,(\pi)$. If $\mathrm{dis}\,(\pi) < \epsilon_{\min}$, set $\epsilon_{\min} = \mathrm{dis}\,(\pi)$ and $\pi_{\min} = \pi$.

3.2. Local refinement

The second stage of our approach is a local refinement. The greedy algorithm followed by the optional branch and bound global minimization finds the minimum distortion non-trivial mapping $\pi_{\min} : X_R \to X_R$. We use $\pi_{\min}$ as a coarse-grid initialization for a multiresolution optimization scheme used to solve the GMDS problem.

We optimize over the images $x'_i = g(x_i)$,

$$\min_{\{x'_1,...,x'_N\} \subset \mathbb{F}(\hat{X})} \max_{i,j=1,...,N} \left|\hat{d}_{ij} - \hat{d}_X(x'_i, x'_j)\right|, \qquad (8)$$

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 [4]. The local minimizer of (8), denoted by $\{x^*_1, ..., x^*_N\}$, is found by means of convex optimization. A guaranteed convergence to a non-trivial solution (i.e., a solution distinct from the identity map) can be obtained by constraining the maximum distance in $\mathbb{F}(\hat{X})$ between $\pi_{\min}$ and g in the GMDS problem.

The global asymmetry is computed as

$$\mathrm{asym}(\hat{X}) \quad = \quad \max_{i,j=1,...,N} \left|\hat{d}_{ij} - \hat{d}_X(x^*_i, x^*_j)\right|,$$

and the local asymmetry as

$$\mathrm{asym}(\hat{X}, x_i) \quad = \quad \max_{j=1,...,N} \left|\hat{d}_{ij} - \hat{d}_X(x^*_j, x_i)\right|.$$

4. Results

4.1. Symmetry detection

In order to assess the accuracy of our method, we performed an experiment of intrinsic symmetry detection on a data set of four different non-rigid shapes (dog, giraffe, man and crocodile). Each shape appeared in three instances: extrinsically symmetric (serving as a reference), intrinsically symmetric but not extrinsically symmetric (obtained by means of a near-isometric deformation of the reference shapes) and asymmetric (obtained by local non-isometric deformation). The local asymmetric features included enlarged ear of the dog, elongated leg of the giraffe, cropped hand of the man and enlarged foot of the crocodile.

Figure 2 visualizes our intrinsic symmetry detection method. The color represents the local measure of asymmetry ($\mathrm{asym}(\hat{X}, x_i)$). The numbers represent the global measure of asymmetry ($\mathrm{asym}(\hat{X})$). In most cases, we were able to identify symmetric shapes and correctly detect the symmetry breaking features (marked with arrows in Figure 2).

4.2. Visualization of the space of symmetries

Extrinsic symmetries of a shape form a sub-group of the isometry group of $\mathbb{R}^3$, and thus its members can be easily parametrized (for example, in case of a rotational symmetry, such parameters can be a unit direction vector and an angle). This allows a simple description of the extrinsic symmetry group. Unfortunately, such a straightforward parametrization is generally unavailable for intrinsic symmetries, and even less for their approximate counterparts. Information about the intrinsic symmetries of a shape can be inferred from the structure of the space $\mathbb{F}(X)$ and the associated distortion $\mathrm{dis} : \mathbb{F}(X) \to \mathbb{R}$. Despite the very high dimensionality of $\mathbb{F}(X)$, its metric structure can be approximately visualized the as a configuration of points in the Euclidean space, where each point represents a map $g \in \mathbb{F}(X)$, while the Euclidean distance between each pair of points approximates $d_{\mathbb{F}(X)}$. Such a representation can be straightforwardly created using MDS techniques.

An approximation of the distortion function is obtained by projecting the values of $\mathrm{dis}(g)$ onto corresponding points in the representation space. The level sets of the approximate distortion function reveal the structure of the intrinsic symmetry group, exhibiting a pattern of local minima. Clearly pronounced local minima correspond each to a different symmetry.

Figure 3 visualizes the structure of the approximate symmetries space $\mathbb{F}(X)$ of a shape obtained by isometrically folding a planar patch. Such a shape has eight intrinsic symmetries which are described by a dihedral group D_4 (as opposed to only one extrinsic symmetry obtained by reflection with respect to a diagonal). Eight local minima, corresponding to the dihedral symmetries can be clearly distinguished.

5. Conclusions

We formulated the problem of detecting symmetries of non-rigid shapes. Our measure of symmetry relies on the intrinsic geometry of the shape, which allows to find symmetries that are insensitive to bending and deformation. We presented a practical approach for the numerical computation of intrinsic symmetries based on the GMDS method. Experimental results demonstrate the accuracy of our approach.

References

[1] H. Alt, K. Mehlhorn, H. Wagener, and E. Welzl. Congruence, similarity, and symmetries of geometric objects. *Discrete Comput. Geom.*, 3:237–256, 1988. 2

[2] M. J. Atallah. On symmetry detection. *IEEE Trans. Computers*, c-34(7), July 1985. 2

[3] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103:1168–1172. 2, 3

[4] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, January 2006. 4

[5] K. Cheung and H. Ip. Symmetry detection using complex moments. *International Conference on Pattern Recognition (ICPR)*, 2:1473–1475, 1998. 2

[6] H. Cornelius and G. Loy. Detecting rotational symmetry under affine projection. *International Conference on Pattern Recognition (ICPR)*, 2:292–295, 2006. 2

[7] 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. 2

[8] A. Elad and R. Kimmel. Bending invariant representations for surfaces. In *Proc. CVPR*, pages 168–174, 2001. 2, 3

[9] Y. Eldar, M. Lindenbaum, M. Porat, and Y. Y. Zeevi. The farthest point strategy for progressive image sampling. *IEEE Trans. Image Processing*, pages 1305–1315, 1997. 3

[10] N. Gelfand, N. Mitra, L. Guibas, and H. Pottmann. Robust global registration. *SGP'05*, 2005. 4

[11] Y. Gofman and N. Kiryati. Detecting symmetry in grey level images: The global optimization approach. *International Conference of Pattern Recognition (ICPR)*, 1996. 2

[12] S. Haraguchi, K. Takada, and Y. Yasuda. Facial Asymmetry in Subjects with Skeletal Class III Deformity. *Angle Orthodontist*, 72(1):28–35. 2

[13] C. Huisinga-Fischer, J. Souren, F. Werken, B. Prahl-Andersen, and F. van Ginkel. Perception of Symmetry in the Face. *Journal of Craniofacial Surgery*, 15(1):128–134, 2004. 1

[14] M. Kazhdan, B. Chazelle, D. Dobkin, T. Funkhouser, and S. Rusinkiewicz. A reflective symmetry descriptor for 3D models. *Algorithmica*, 38(1):201–225, 2003. 2

[15] J. Kepler. *Strena seu de nive sexangula*. Frankfurt, 1611. 1

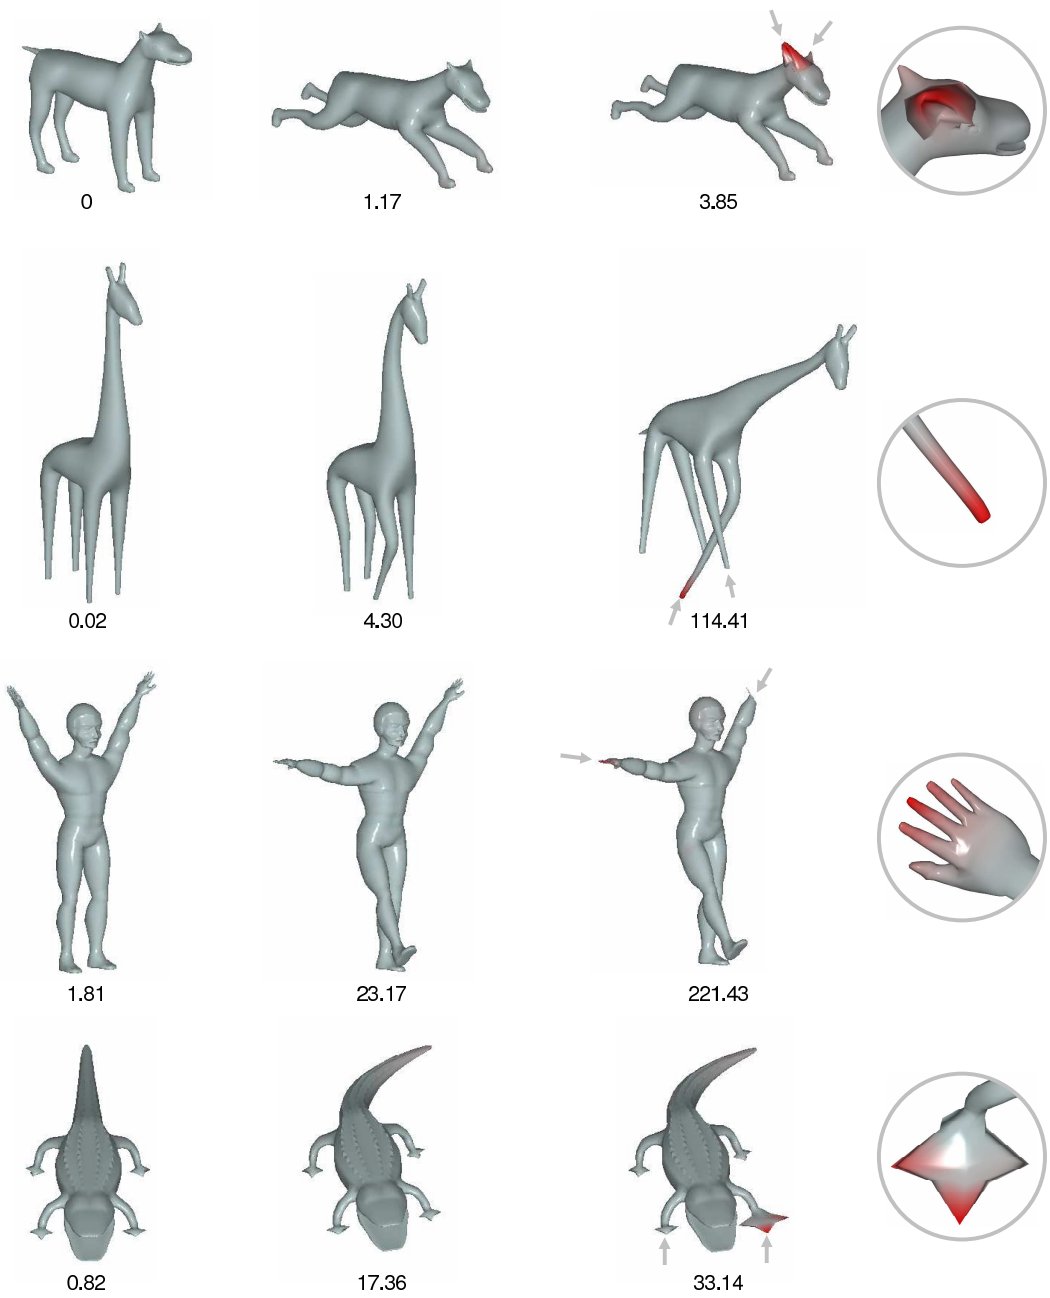

Figure 2. Detection of symmetries of non-rigid shapes. Color encodes the local measure of asymmetry (red correspond to large local asymmetry). Numbers represent the global asymmetry measure (larger means less symmetric).

[16] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proc. National Academy of Sciences*, 95(15):8431–8435, 1998. 3

[17] G. Loy and J. Eklundth. Detecting symmetry and symmetric constellations of features. *Proc. ECCV*, 2:508–521, 2006. 2

[18] M. Mancas, B. Gosselin, and B. Macq. Fast and automatic tumoral area localisation using symmetry. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2:725–728, 2005. 1

[19] G. Marola. On the detection of axes of symmetry of symmetric and almost symmetric planner images. *IEEE Trans. PAMI*, 11(1), Jan. 1989. 2

[20] L. Mealey, R. Bridgstock, and G. Townsend. Symmetry and perceived facial attractiveness: a monozygotic co-twin comparison. *Journal of Personality and Social Psychology*, 76(1):151–158, 1999. 1

[21] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud

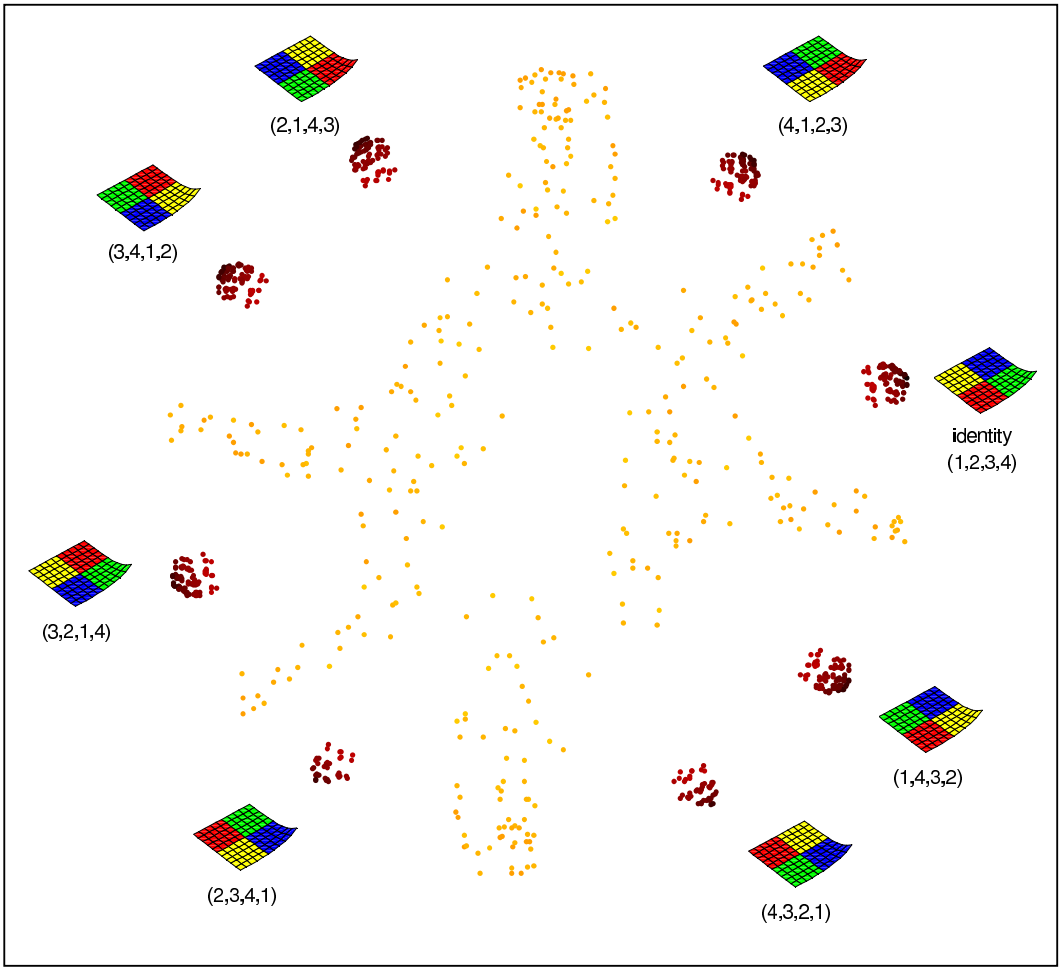

Figure 3. Visualization of the space $\mathbb{F}(X)$ of a folded planar patch. Each point corresponds to a member of $\mathbb{F}(X)$ (mapping from X to itself). Euclidean distances between the points represent $d_{\mathbb{F}(X)}$. Color represents the distortion dis (darker implies smaller distortion). The dihedral symmetry is depicted by showing the permutations of the corners of the shape.

data. *Foundations of Computational Mathematics*, 5:313–346, 2005. 2, 3

[22] N. J. Mitra, L. J. Guibas, and M. Pauly. Partial and approximate symmetry detection for 3D geometry. *SIGGRAPH*, pages 560–568, 2006. 2

[23] C. Moenning and N. Dodgson. A new point cloud simplification algorithm. *International Conference on Visualization, Imaging and Image Processing*, 2003. 3

[24] F. D. Natale, D. Giusto, and F. Maccioni. A symmetry-based approach to facial features extraction. *International Conference on Digital Signal Processing Proceedings (ICDSP)*, 2:521–525, 1997. 2

[25] G. Peyré and L. Cohen. Geodesic Remeshing Using Front Propagation. *International Journal of Computer Vision*, 69(1):145–156, 2006. 3

[26] D. Reisfeld and Y. Yeshurun. Robust detection of facial features by generalized symmetry. *International Conference on Pattern Recognition (ICPR)*, 1:117–120, 1992. 2

[27] I. Shimshoni, Y. Moses, and M. Lindernbaumplr. Shape reconstruction of 3D bilaterally symmetric surfaces. *International Conference on Image Analysis and Processing (ICIP)*, page 76, 1999. 2

[28] C. Sun and J. Sherrah. 3D symmetry detection using the extended gaussian image. *IEEE Trans. PAMI*, 19(2):164–168, 1997. 2

[29] H. Weyl. *Symmetry*. Princeton University Press, 1983. 1

[30] J. Wolter, T. Woo, and R. Volz. Optimal algorithms for symmetry detection in two and three dimensions. *The Visual Computer*, 1:37–48, 1985. 2

[31] H. Zabrodsky, S. Peleg, and D. Avnir. Symmetry as a continuous feature. *IEEE Trans. PAMI*, 17(12):1154–1166, 1995. 2