Eurographics Workshop on 3D Object Retrieval (2012)
M. Spagnuolo, M. Bronstein, A. Bronstein, and A. Ferreira (Editors)

Intrinsic Local Symmetries: A Computational Framework

Carmi Grushko Dan Raviv Ron Kimmel

Dept. of Computer Science
Technion, Israel
{carmi,darav,ron}@cs.technion.ac.il

Abstract

We present a computational framework for finding metric-preserving tangent vector fields on surfaces, also known as Killing Vector Fields. Flows of such vector fields define self-isometries of the surface, or in other words, symmetries. Our approach is based on general-purpose isometry-finding frameworks, and is shown to be robust to noise. In addition, we demonstrate symmetry recovery using non-Euclidean metrics.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—

1. Introduction

The symmetry of a surface can be defined in terms of self-similarity: a cupboard cup has rotational symmetry because it can be rotated about one axis with no apparent change. Squash the cup, and the symmetry would be lost - or would it? From the point of view of an infinitesimally small ant on its surface, squashing the cup has no effect - if all points on the creased cup were to move as if the cup hasn't been altered, our ant would not be able to tell the difference made by such a deformation.

We define self-similarity in terms of intrinsic properties of surfaces. Such symmetry is said to be *intrinsic*, because it does not depend on properties of the embedding space. Rotational symmetry, such as the one a cup possesses, is an example of an *extrinsic* symmetry, one which is defined in terms of transformations in the embedding space.

Such rotational symmetry is also *infinitesimal,* since we can apply it in arbitrarily small amounts. By contrast, the reflective symmetry of Leonardo da Vinci's Vitruvian Man is not.

Infinitesimal symmetries of surfaces can be described by a tangent vector field; once obtained, such a field can be used as a shape descriptor. Possible descriptors include a histogram of norms of the vectors, their hodograph, and methods inspired by descriptors for color images.

We present a framework for explicitly computing intrinsic infinitesimal symmetries, using Generalized Multidimensional Scaling [BBK06b], and demonstrate its robustness to noise.

1.1. Mathematical Background

The notion of distance measurement on manifolds, and in particular on surfaces, is formalized through the use of metric tensors. A metric tensor at point p of manifold M maps two tangent vectors at this point to a scalar,

$$g_p : T_pM \times T_pM \to \mathbb{R}^+ \cup \{0\}, \quad (1)$$

and essentially defines an inner-product on tangent vectors. A manifold equipped with a bilinear, semi-positive definite, symmetric and non-degenerate metric tensor is called a *Riemannian Manifold*. Consequently, the angle between two tangent vectors u and v is defined through $g_p(u,v)$, and more relevant to us, the size of a tangent vector $v \in T_pM$ is defined as $\sqrt{g_p(v,v)}$. Given a curve on the manifold, we may compute its length by integrating the metric over its trace. Furthermore, if the manifold is bounded, connected and complete, for every two points there exists a shortest path connecting them, called a *minimal geodesic*. We may therefore induce a metric $d(p,q)$ on the manifold, where the distance between two points is defined as the length of the shortest path between them.

A self-isometry of a manifold is a map $\phi : M \to M$ which

DOI: 10.2312/3DOR/3DOR12/033-038

preserves inter-point distances, that is,

$$\forall p,q \in M \quad d(p,q) = d(\phi(p),\phi(q)), \tag{2}$$

and essentially captures the notion of intrinsic symmetry of a manifold. We may also define ε-isometries, where the above equality is replaced by a neighborhood of size ε.

A map $\phi : M \to M$ is *infinitesimal* if there exists a family of maps $\varphi^t : M \to M$ parametrized by $t \in \mathbb{R}$, such that for all $\delta > 0$, there exists a set $T_\delta = \{t_i\}_{i=1}^n \subset (0,\delta)$ of positive numbers smaller than δ, for which we have

$$\phi \equiv \varphi^{t_1} \circ \varphi^{t_2} \circ \cdots \circ \varphi^{t_n}. \tag{3}$$

Intuitively, an infinitesimal map is a continuous flow of points on the manifold.

A continuous, infinitesimal map can be described by a continuous tangent vector field $U : M \to T_pM$. In this case the self-isometry becomes

$$\phi(p) = \exp_p(U(p)), \tag{4}$$

where $\exp_p(\cdot)$ is the exponential map at point p, and is termed the *flow* of the vector field.

A tangent vector field which preserves the metric tensor of a manifold also preserves inter-point distances, and is called a *Killing Vector Field* (KVF). Although KVFs are rare [Mye36], approximate KVFs (AKVFs), which are near-isometric vector fields, occur naturally almost everywhere. For example, socks are intuitively symmetric, but do not posses an infinitesimal symmetry from a pure mathematical point of view.

1.2. Previous work

Traditionally, symmetry extraction methods were based on exploiting the structure of an embedding space of a surface, [BBW*09, PMW*08, SS97]; such methods are inappropriate where isometric deformations are allowed, such as bending. Methods that find intrinsic symmetries exist, [RBBK07, OSG08, XZT*09], but they focus on discrete symmetries.

In [Mat68], the problem of finding AKVFs is stated as a functional optimization problem which can be transformed into an eigenvalue problem. An eigenvector which corresponds to the zero eigenvalue is a true KVF, and the larger the eigenvalue, the farther its corresponding eigenvector is from being metric-preserving. An implementation of a similar method was the focus of a paper by Ben-Chen et al.[BCBSG10]. It performs well on smooth surfaces, but is based on local (differential) properties of the surface, and is therefore sensitive to noise as we show later in this paper.

An alternative approach to finding AKVFs is to find a general self-isometry of a surface, restricting it to be continuous and infinitesimal. A series of papers by Bronstein et al. (e.g.,[BBK06a, BBK06b, BBK*10]) focused on finding isometries between surfaces. Their proposed algorithm, *GMDS*, is a natural candidate for finding AKVFs.

1.3. Generalized MDS

Loosely speaking, GMDS [BBK06b] finds a correspondence between two sets of points that reside on triangulated surfaces. A byproduct of this correspondence is a quantitative dissimilarity measure between those surfaces, which can be used for example for face recognition.

Let S and Q be two Riemannian manifolds of dimension two with metrics d_S and d_Q respectively, and let $X = \{x_i\}_{i=1}^N \subseteq S$ be points in S. Also, define $Y = \{y_i\}_{i=1}^N \subseteq Q$, and require that $N = |X| = |Y| < \infty$ where $|X|$ denotes the number of points in X. GMDS is an iterative algorithm that minimizes the *generalized p-stress* as a function of Y, which for $p < \infty$ is defined as

$$\sigma(Y) = \sum_{i=1}^{N}\sum_{j=1}^{N} \left|d_S(x_i,x_j) - d_Q(y_i,y_j)\right|^p, \tag{5}$$

while for the $p = \infty$ case it is defined as

$$\sigma(Y) = \max_{1\le i,j\le N} \left|d_S(x_i,x_j) - d_Q(y_i,y_j)\right|. \tag{6}$$

Note that throughout GMDS, only Y changes.

Ideally, X and Y are sufficiently dense nets of S and Q respectively. If a global solution is found, the value of $\sigma(Y)$ can qualitatively show how isometric S and Q are: a zero value obviously tells us that the manifolds are isometric. Conversely, a large value means the manifolds are intrinsically different. See [MS05] for an analysis in the context of surface sampling and the Gromov–Hausdorff distance.

2. Implementation

As stated before, we are interested in computing a non-trivial vector field $U : M \to T_pM$ which preserves inter-point distances; that is, given a set of points $P \subset M$ on the surface and their inter-point distances matrix $D(P)$, we would like the following to hold: $D(\exp_P(U(P))) \approx D(P)$.

A discrete approximation to this problem would be to find two different sets of points $X = \{x_i\}_{i=1}^N$ and $Y = \{y_i\}_{i=1}^N$ on a triangulation of the surface M which minimize the *p-norm stress functional*

$$\arg\min_{i<j} \left\|d(x_i,x_j) - d(y_i,y_j)\right\|_p, \tag{7}$$

where $d : M^2 \to \mathbb{R}^+ \cup \{0\}$ is a metric on M, and most importantly, X is close to Y, or in other words, $d(x_i,y_i)$ is small for all $i \in [N]$. Equivalently, one may define $\sigma_{ij} = \left|d(x_i,x_j) - d(y_i,y_j)\right|$ as the *stress between point i and point j*, and minimize the elements of this stress matrix.

2.1. Using GMDS

In practice, we initialize X using the vertices of a triangulation of M, restrict one point to a new location, and then translate the rest of the points to new locations on M which minimize (7). Note that restricting a single point to a *new* location is essential in order to avoid the trivial solution to the problem, which corresponds to the null vector field. It is essentially a specific case of GMDS, where $S = Q$. The metric $d(x,y)$ can be arbitrary and only has to be defined on the vertices, and values at other points are linearly interpolated according to [BBK09] (pp 197-198). This means the metric is only C^0 continuous, which forces points to only move within their own triangles per step - moving across the mesh therefore requires multiple steps. Moreover, because moving all the points at once is a non-convex problem and is hard to solve, we move one point at a time, freezing the rest of the points. Since the mesh is composed of triangles, we represent points using their barycentric coordinates $u_1, u_2, 1-u_1-u_2$ and their triangle number t.

We choose the $p = \infty$ norm in the context of the stress minimization problem (7), because this norm is very sensitive to the relocation of a single point. This allows us to displace a single point, which forces all other points to move. Using the L_2 norm, for example, might lead to a situation where most of the stress is concentrated around one point, which corresponds to almost no movement of the rest of the points. The resulting optimization problem we solve for each point we move is therefore

$$\arg_{u_1^k,u_2^k} \min_{i\neq k} \max \left| d(x_i,x_k) - d(y_i,q) \right|, \tag{8}$$

where $q \triangleq \left[t^k, u_1^k, u_2^k, 1-u_1^k-u_2^k\right]$. This can be re-written as the following constraint optimization problem over the variables u_1, u_2 and ε (note that t^k does not change)

$$\begin{aligned} \min \quad & \varepsilon \\ \text{s.t.} \quad & \left| d(x_i,x_k) - d(y_i,q) \right| \leq \varepsilon. \end{aligned} \tag{9}$$

As the distance matrix $\left\{d(x_i,x_j)\right\}_{i,j}^N$ does not change, we denote it as δ_{ij}. Because $d\left(y_i, \left[t^k, u_1^k, u_2^k, 1-u_1^k-u_2^k\right]\right)$ is linear when t^k does not change, we rewrite it as $C^{ik} \cdot (u_1,u_2,1)$ and obtain the following non-negative linear optimization problem in three variables, u_1, u_2 and ε, and $2(N-1)$ constraints :

$$\begin{aligned} \min \quad & \varepsilon \\ \text{s.t.} \quad & \forall i \neq k : -\varepsilon \leq \delta_{ik} - C^{ik} \cdot (u_1,u_2,1) \leq \varepsilon \\ & u_1 + u_2 \leq 1 \\ & u_1, u_2 \in [0,1]. \end{aligned} \tag{10}$$

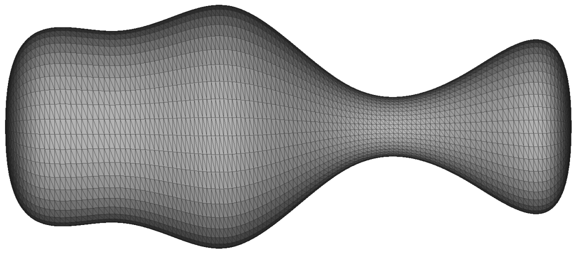

Figure 2: *A Bowling-pin Model*

3. Experimental Results

3.1. Symmetry finding

In order to evaluate the performance of our proposed algorithm, several models were tested. The results appear in the Figure 1. The restricted point is shown in red. We used geodesic distances, computed using Fast Marching Method [KS98], as the metric $d(x,y)$ in this set of experiments.

3.2. Robustness

We demonstrate the algorithm's resilience to noise using a surface of revolution (Figure 2). We gradually add Gaussian noise in the normal direction and measure the quality of the result in the following way. The resulting vector field is converted into pairs of starting points and ending points, which are then projected to the clean surface. Then, a vector field is reconstructed, and is compared to a reference vector field. We also compare the field to the result of the algorithm by Ben-Chen et al. [BCBSG10], designated as "AKVF".

Vector fields are compared in two ways: the first ignores the length of the vector, and is the sum of (positive) angles between each two vectors which emanate from the same point on the surface. The second way is the Euclidean distance between end-points of each two corresponding vectors.

The results are detailed in Figures 3 and 4. Figure 5 shows a zoom in of the resulting vector field for Gaussian noise with standard deviation of 0.08.

3.3. Different Metrics

There is little limitation as to what metric is used in GMDS. In the following experiment, we use an equi-affine invariant metric [RBB*11] to restore symmetry which is lost due to a compression of the surface. This metric is given in local coordinates as

$$G = \tilde{G} \cdot \left\|\tilde{G}\right\|^{-1/4}, \tag{11}$$

where

$$\tilde{G}_{ij} = \det\left(S_u, S_v, S_{ij}\right) \tag{12}$$

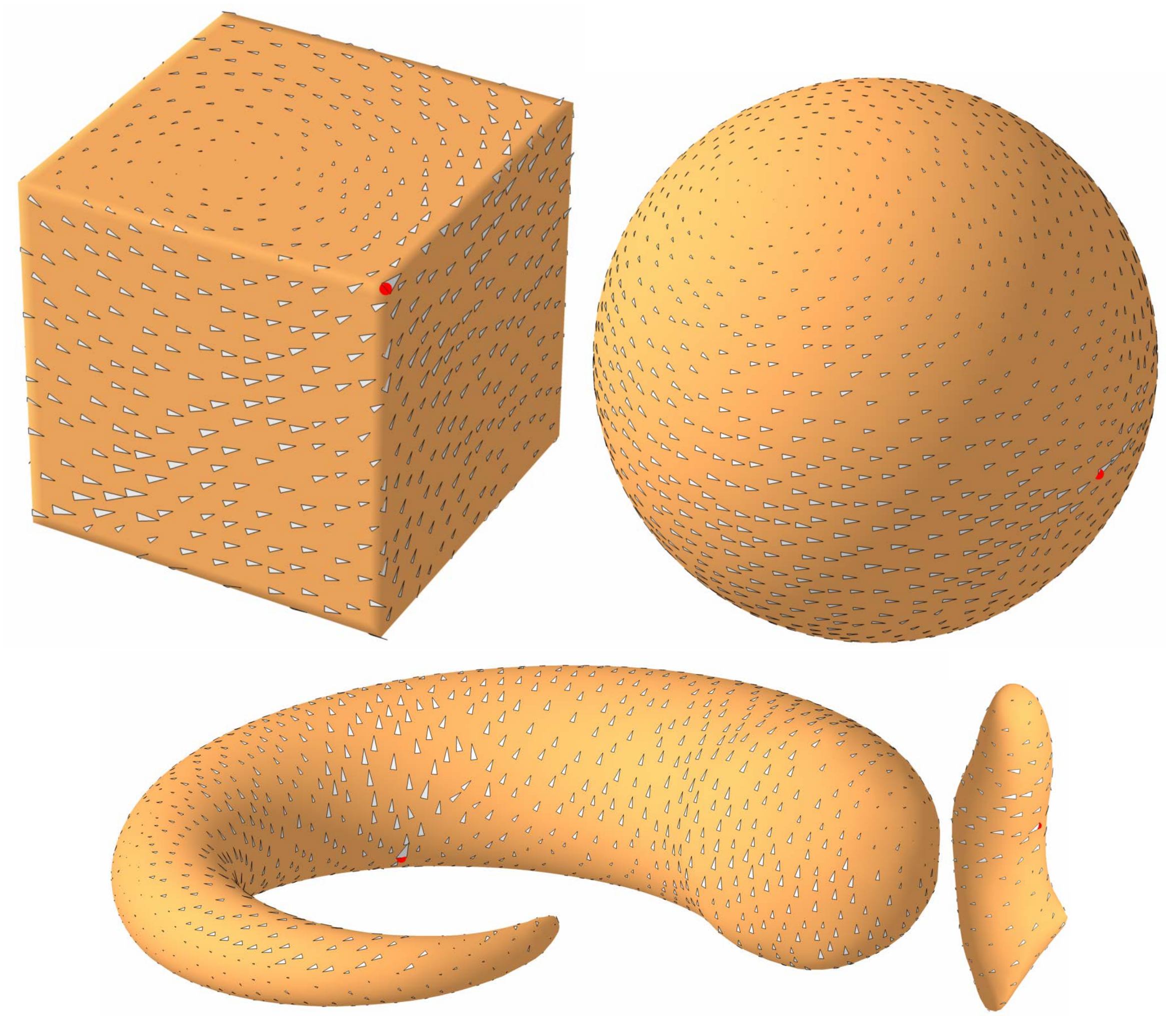

Figure 1: *The proposed method applied to various surfaces.*

and S_{ij} are the second derivatives of the surface S, such that $S_{12} \triangleq S_{uv}$, $S_{22} \triangleq S_{vv}$ and $S_{11} \triangleq S_{uu}$. Because $\tilde{G}$ is not positive-definite at points with negative Gaussian curvature, it is eigen-decomposed into VDV^T and $\tilde{G}$ is redefined as

$$\tilde{G} = V \cdot \text{abs}(D) \cdot V^T. \tag{13}$$

We begin with a sphere centered at the origin, which has three linearly-independent KVFs. We then multiply all its x coordinates by some constant, resulting in an ellipsoid. See Figure 6. This deformation has left but a single KVF, which coincides with the rotational symmetry of the shape.

We ran our algorithm on the ellipsoid, initializing GMDS to look for a KVF along a meridian, once with a geodesic-distance metric, and once with the equi-affine invariant metric. In order to emphasize the results, the ellipsoid was re-inflated back to a sphere. Figure 7 shows the results.

4. Conclusions

We introduced an algorithm for computing infinitesimal intrinsic symmetries of surfaces using a general-purpose intrinsic isometry solver. We also showed this algorithm is robust to noise, and that it is applicable to use other, non geodesic-distance metrics to compute such symmetries.

5. Acknowledgments

The authors would like to thank Dr. Mirela Ben-Chen for sharing models and source code for finding AKVFs.

This research was supported by European Community's FP7-ERC program, grant agreement no. 267414.

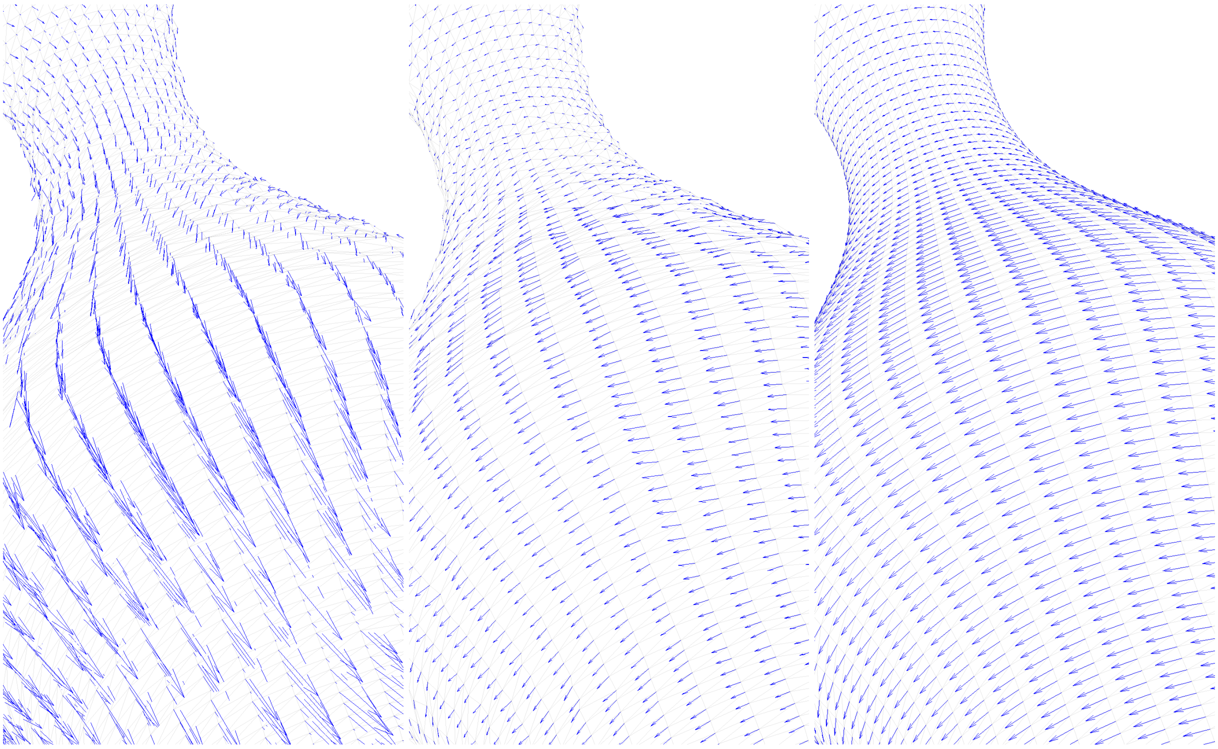

Figure 5: *Left: AKVF. Middle: GMDS. Right: Reference vector field*

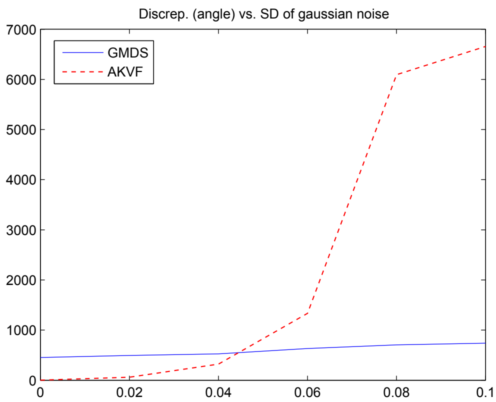

Figure 3: *Quality vs. Standard Deviation (angle discreprency)*

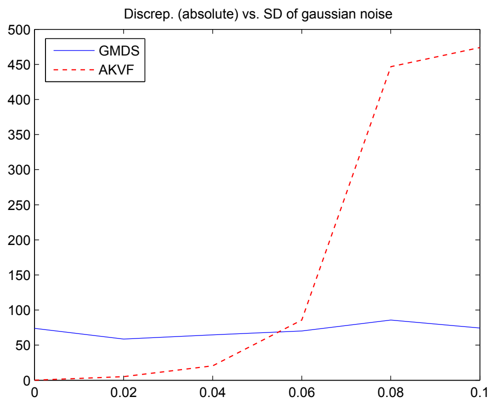

Figure 4: *Quality vs. Standard Deviation (absolute discreprency)*

References

[BBK06a] Bronstein A. M., Bronstein M. M., Kimmel R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Scientific Computing 28*, 5 (2006), 1812–1836. 2

[BBK06b] Bronstein A. M., Bronstein M. M., Kimmel R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. National Academy of Sciences (PNAS) Volume 103/5* (January 2006), 1168–1172. 1, 2

[BBK09] Bronstein A. M., Bronstein M. M., Kimmel R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2009. 3

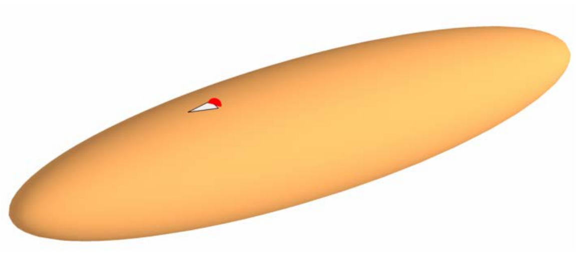

Figure 6: *Ellipsoid resulting from compressing a sphere. Arrow shows the direction in which a KVF is searched for.*

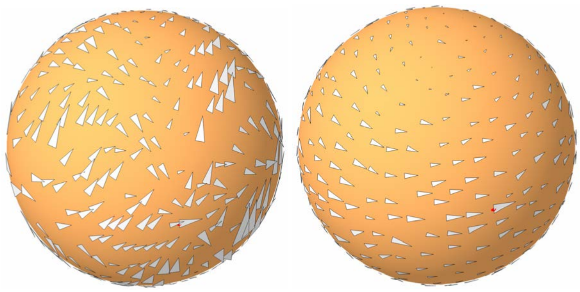

Figure 7: *Re-inflated ellipsoid. Left: Geodesic distances. Right: Equi-affine invariant metric.*

[BBK*10] BRONSTEIN A., BRONSTEIN M., KIMMEL R., MAHMOUDI M., SAPIRO G.: A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision 89*, 2 (2010), 266–286. 2

[BBW*09] BOKELOH M., BERNER A., WAND M., SEIDEL H., SCHILLING A.: Symmetry detection using line features. In *Computer Graphics Forum (Special Issue of Eurographics)* (2009), vol. 28, pp. 697–706. 2

[BCBSG10] BEN-CHEN M., BUTSCHER A., SOLOMON J., GUIBAS L. J.: On discrete killing vector fields and patterns on surfaces. *Comput. Graph. Forum 29*, 5 (2010), 1701–1711. 2, 3

[KS98] KIMMEL R., SETHIAN J. A.: Computing geodesic paths on manifolds. In *Proc. Natl. Acad. Sci. USA* (1998), pp. 8431–8435. 3

[Mat68] MATZNER R.: Almost symmetric spaces and gravitational radiation. *Journal of Mathematical Physics 9* (1968), 1657. 2

[MS05] MEMOLI F., SAPIRO G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics 5/3* (2005), 313–347. 2

[Mye36] MYERS S.: Isometries of 2-dimensional Riemannian manifolds into themselves. *Proceedings of the National Academy of Sciences of the United States of America 22*, 5 (1936), 297. 2

[OSG08] OVSJANIKOV M., SUN J., GUIBAS L.: Global intrinsic symmetries of shapes. In *Computer graphics forum* (2008), vol. 27, Wiley Online Library, pp. 1341–1348. 2

[PMW*08] PAULY M., MITRA N., WALLNER J., POTTMANN H., GUIBAS L.: Discovering structural regularity in 3D geometry. *ACM Transactions on Graphics (TOG) 27*, 3 (2008), 1–11. 2

[RBB*11] RAVIV D., BRONSTEIN M. M., BRONSTEIN A. M., KIMMEL R., SOCHEN N. A.: Affine-invariant diffusion geometry for the analysis of deformable 3d shapes. In *CVPR* (2011), pp. 2361–2367. 3

[RBBK07] RAVIV D., BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Symmetries of non-rigid shapes. In *ICCV* (2007), IEEE, pp. 1–7. 2

[SS97] SUN C., SHERRAH J.: 3D symmetry detection using the extended gaussian image. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 19*, 2 (1997), 164–168. 2

[XZT*09] XU K., ZHANG H., TAGLIASACCHI A., LIU L., LI G., MENG M., XIONG Y.: Partial intrinsic reflectional symmetry of 3D shapes. *ACM Transactions on Graphics (TOG) 28*, 5 (2009), 138. 2