Properly-constrained Orthonormal Functional Maps for Intrinsic Symmetries

Xiuping Liu[a], Shuhua Li[a], Risheng Liu[b], Jun Wang[c], Hui Wang[d,*], Junjie Cao[a,e,**]

[a] School of Mathematical Sciences, Dalian University of Technology, China
[b] School of Software Technology, Dalian University of Technology, China
[c] College of Mechanical and Electrical Engineering, Nanjing University of Aeronautics and Astronautics, China
[d] School of Information Science and Technology, Shijiazhuang Tiedao University, China
[e] School of Mathematics and Information Sciences, Nanchang Hangkong University, China

Abstract

Intrinsic symmetry detection, phrased as finding intrinsic self-isometries, courts much attention in recent years. However, extracting dense global symmetry from the shape undergoing moderate non-isometric deformations is still a challenge to the state-of-the-art methods. To tackle this problem, we develop an automatic and robust global intrinsic symmetry detector based on functional maps. The main challenges of applying functional maps lie in how to amend the previous numerical solution scheme and construct reliable and enough constraints. We address the first challenge by formulating the symmetry detection problem as an objective function with descriptor, regional and orthogonality constraints and solving it directly. Compared with refining the functional map by a post-processing, our approach does not break existing constraints and generates more confident results without sacrificing efficiency. To conquer the second challenge, we extract a sparse and stable symmetry-invariant point set from shape extremities and establish symmetry electors based on the transformation, which is constrained by the symmetric point pairs from the set. These electors further cast votes on candidate point pairs to extract more symmetric point pairs. The final functional map is generated with regional constraints constructed from the above point pairs. Experimental results on TOSCA and SCAPE Benchmarks show that our method is superior to the state-of-the-art methods.

Keywords: global intrinsic symmetry, functional map, moderate non-isometric deformation

1. Introduction

Symmetry is an universal phenomenon in nature which provides global information about the structure of objects. Numerous geometry processing tasks, such as shape matching [1], segmentation [2], geometry completion [3] and meshing [4], benefit from symmetry information. Hence a great deal of work [5] devotes to extract symmetries from geometric data, *e.g.*, point clouds data and polygon meshes.

Most of the previous work concentrates on extrinsic symmetries [6, 7]. Recently, intrinsic symmetry detection, phrased as finding intrinsic self-isometries, has received more attention, since intrinsic symmetric objects or phenomenons are more common in real world, such as a human in different poses. However, it is infeasible to search the space of non-rigid transformations directly in classical point-to-point representation. So many methods limit the search space to a set of feature points, and adopt combinatorial algorithms to prune point pairs without preserving local geometric similarity and distance structure [8], which are computationally expensive and sensitive to geodesic noises. Kim et al. [9] take advantage of the fact that intrinsic self-isometries are contained in a low dimensional Möbius transformation space [10] to select the best self-isometry. The symmetry-invariant set, used to generate candidate Möbius transformations, consists of some local extrema of the Average Geodesic Distance function (AGD) [11]. The set may be not perfectly symmetric and leads to failure results. Ovsjanikov et al. [12] extract intrinsic symmetries using functional maps [13]. But they need at least one reference shape with a known symmetry to estimate the quotient space and a consistent decomposition to obtain the final dense intrinsic symmetries. The decomposition divides the shape into fundamental domains, *e.g.*, the right part and the left part of the shape in the case of reflectional symmetry. Furthermore, shapes undergoing considerable degree of non-isometric deformations, such as humanoid models with connections between torso and

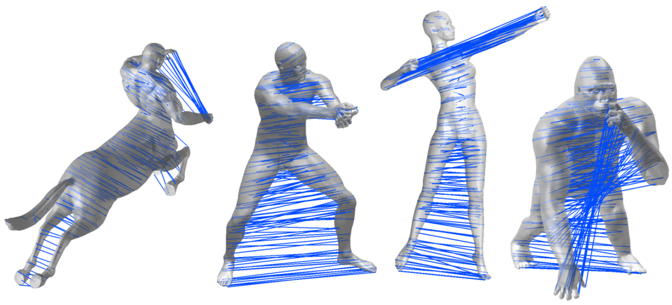

Figure 1: The results of our method for nearly self-isometric shapes (centaur, michael, victoria and gorilla).

*Corresponding author. E-mail: wangh@stdu.edu.cn (H. Wang). Telephone: 86-15076171760.
**Corresponding author. E-mail: jjcao@dlut.edu.cn (J. Cao).

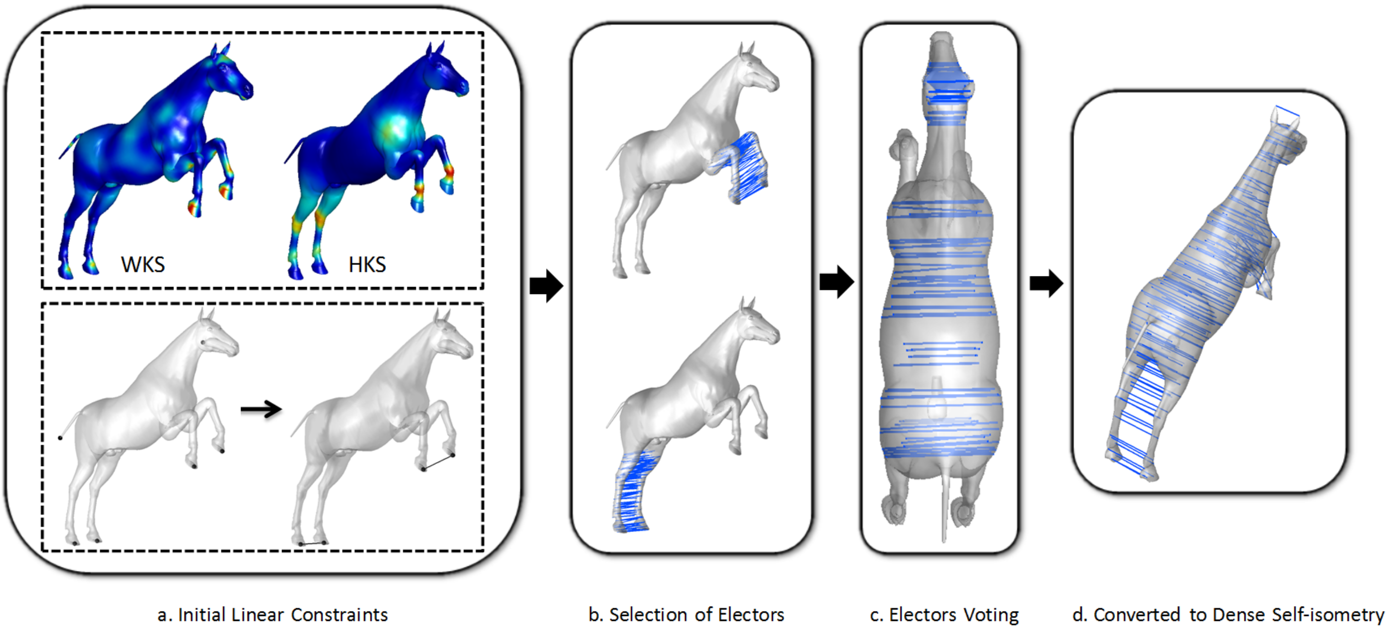

Figure 2: The pipeline of our method.

other parts, also challenge the existing methods.

We observe that most existing methods detect intrinsic symmetry over a sparse set of feature points, then propagate the sparse correspondence to the entire shape using geodesic distance. The performance is degenerated since the propagation only considers metric. The functional map framework presents a compact representation of correspondences between shapes, and provides an efficient way to convert functional maps into dense point-to-point correspondences [13]. This motivates us to present an automatic and robust method for global intrinsic symmetry detection leveraging the functional map representation (Fig. 1). Intrinsic symmetries are non-trivial self-isometries represented by orthonormal functional map matrixes. Extending the functional map to detect global intrinsic symmetry directly suffers from the absence of constraints indicating the underlying non-trivial self-isometry. Existing descriptors, such as Heat Kernel Signature (HKS) [14] and Wave Kernel Signature (WKS) [15], provide no valuable cues for distinguishing identity transformation with other symmetry transformations, since they remain invariant in these transformations. Point or segment correspondences contain useful information for distinguishing the above transformations, however the establishment of reliable and enough symmetric point or segment pairs itself is a challenge problem. The key idea of our method is to construct reliable and sufficient regional constraints from symmetric point pairs. The most prominent and stable feature pairs tend to lie on the extremities of the model. We design an initialization procedure to extract sparse and reliable symmetric point pairs from the extremities, and a voting procedure to extract more symmetric point pairs.

In the initialization procedure, initial symmetric point pairs are chosen from a symmetry-invariant set (Fig. 2 (a)), which is extracted from shape extremities and whose stability and sparseness make the procedure reliable and efficient. Then we compute an initial functional map satisfying regional constraints, constructed from the initial point pairs. We specify the parts containing the initial symmetric pairs as the reliable parts of the initial functional map. More symmetric point pairs over the reliable parts are selected as symmetry electors (Fig. 2 (b)). In the following procedure, a voting scheme is proposed to extract more symmetric point pairs outside the reliable parts (Fig. 2 (c)). The final functional map is generated with the regional constraints constructed from all of the point pairs, and converted to a point-to-point mapping (Fig. 2 (d)).

When solving for the functional maps, we formulate the problem as an optimization problem with descriptor, regional and orthogonality constraints simultaneously. Compared with refining the functional map by a post-processing [12, 13, 16], our method does not break other constraints and generates more confident results without sacrificing efficiency. The functional representation, efficient optimization method and effective regional constraints together make our method a faster, automatic and robust implementation for global intrinsic symmetry detection. We demonstrate the effectiveness of our optimization with orthogonality constraints and the voting scheme experimentally (Section 5.1 and Section 5.2). The pipeline of our method is given in Fig. 2. The main contributions are as follows:

1. We present a robust intrinsic symmetry detection method based on functional maps. By formulating the problem as an objective function with descriptor, regional and orthogonality constraints and solving it directly, more faithful results are generated without compromising efficiency.
2. A reliable symmetry-invariant point set is generated by moving a very sparse set of samples towards the extremities of shapes, making the establishment of the initial symmetric point pairs feasible.

3. A voting scheme is proposed to establish more symmetric point pairs, which provides constraints outside the reliable parts.

2. Related work

The vast majority of existing work has been on extrinsic symmetry detection [6, 7, 17, 18, 19]. Recently, intrinsic symmetry detection has attracted more attention [8, 9, 12, 20, 21, 22]. As intrinsic symmetry is a special case of correspondences, most methods work for shape correspondences as well as intrinsic symmetries. Some previous work aims to establish point pairs from feature points [8, 20, 23, 24, 25, 26, 27, 28], which has unknown complexity associated with the number of the feature points. For example, Au et al. [28] prune bad correspondences over skeletal feature nodes. Zhang et al. [27] perform the pruning procedure over prominent shape extremities, which are local extrema of AGD. The local extrema of AGD are unstable under deformations and some local extrema may be not symmetry-invariant points and do not have symmetric points (the bottom row of Fig. 3). Moreover, the number and the location of the local extrema are related to a smoothing parameter, with which AGD is smoothed. Hence, in the pruning step, we establish initial symmetric pairs over a sparse and stable symmetry-invariant set (the top row of Fig. 3), which is extracted from shape extremities. Similar to [28], our pruning-based initialization step is followed by a voting procedure. In the voting step, Au et al. [28] establish electors and candidates over the same set of feature points, and output a sparse correspondence. However, we construct electors and candidates in different parts of shapes. Because more electors and candidates are needed to provide enough regional constraints and solve for a dense intrinsic symmetry via functional maps.

Alternative approaches [8, 20, 21] aim to embed a shape into a new space in which intrinsic symmetry detection is reduced to an extrinsic one. For example, Raviv et al. [8] embed the object into an Euclidean space by generalized multi-dimensional scaling. The original geodesic distances are preserved in the form of corresponding Euclidean distances. They minimize distance distortion directly in the new space. Ovsjanikov et al. [21] define a signature space by the eigenfunctions of the Laplace-Beltrami operator, in which each point is represented as a sequence of signs of the restricted Global Point Signature [29].

Some recent work has attempted to represent intrinsic symmetries as global transformations with a small number of parameters, which is similar to extrinsic symmetry detection [9]. Lipman et al. [10] observe that isometry is a subset of Möbius transformations which has only 6 degrees of freedom for genus zero surfaces, and develop a Möbius Voting scheme to find correspondences of shapes. Kim et al. [9] extend it to detect global intrinsic symmetry since intrinsic symmetries are self-isometries of shapes. Kim et al. [30] blend a large set of candidate conformal maps to form a smooth map, which results in a large blending matrix and is computationally expensive. Liu et al. [31] detect intrinsic reflective symmetry axis curves based on blended intrinsic maps [30]. All of the above methods based on conformal geometry assume the input shapes are genus zero surfaces. The quality of the meshes affects their performance, since they use a mid-edge uniformization technique to map genus zero surfaces onto the extended complex plane.

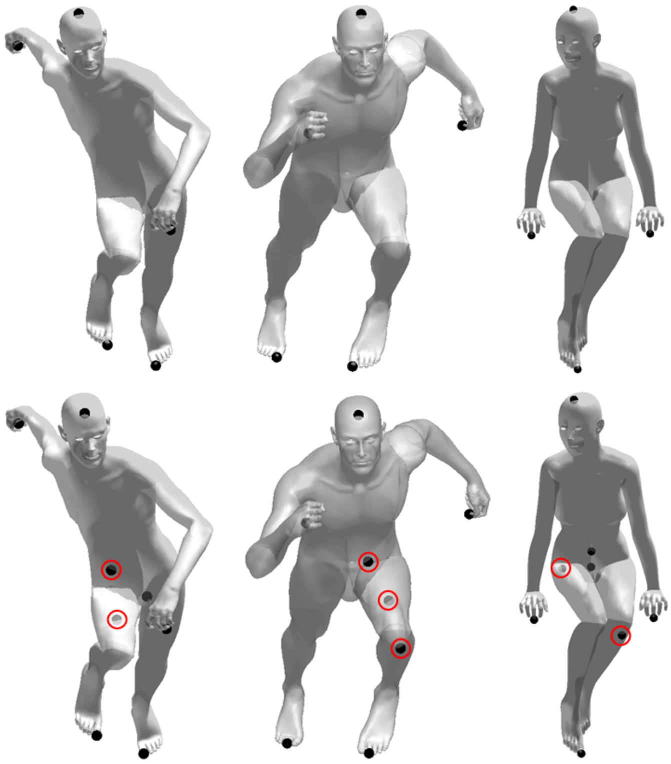

Figure 3: The symmetry-invariant point sets $\mathcal{V}$ (top) and the local extrema of AGD (bottom) on various mesh models. Some local extrema of AGD may be not stationary points and do not have symmetric points (marked by red circle).

Using the novel functional representation [13], Ovsjanikov et al. [12] detect intrinsic symmetry via an appropriate quotient space of the functional space. However, establishing the quotient space requires at least one reference shape with a known symmetry and the conversion to dense intrinsic symmetries is not straightforward.

3. Optimization of global intrinsic symmetry

We adopt functional maps introduced by Ovsjanikov et al. [13] to detect global intrinsic symmetries. Before introducing our objective function in Section 3.2, a brief overview of functional maps is given in Section 3.1.

3.1. Functional maps

Given two compact smooth Riemannian manifolds $\mathcal{M}$, $\mathcal{N}$ and a bijective mapping between them $T : \mathcal{M} \to \mathcal{N}$, a linear transformation between two function spaces is induced $T_F : \mathcal{F}(\mathcal{M}, \mathbb{R}) \to \mathcal{F}(\mathcal{N}, \mathbb{R})$, $T_F(f) = f \circ T^{-1}$, $f \in \mathcal{F}(\mathcal{M}, \mathbb{R})$, where $\mathcal{F}(\mathcal{M}, \mathbb{R})$, $\mathcal{F}(\mathcal{N}, \mathbb{R})$ denote the spaces of real functions on $\mathcal{M}$ and $\mathcal{N}$ respectively.

Given two groups of basis functions, $\{\phi_i^{\mathcal{M}}\}$ of $\mathcal{F}(\mathcal{M}, \mathbb{R})$ and $\{\phi_j^{\mathcal{N}}\}$ of $\mathcal{F}(\mathcal{N}, \mathbb{R})$, the transformation T_F can be fully encoded by a real matrix C defined by

$$T_F(\phi_i^{\mathcal{M}}) = \sum_j c_{ji}\phi_j^{\mathcal{N}}. \quad (1)$$

Conversely, the mapping T can be recovered once the matrix C is obtained, according to **Remark 4.1** in [13]. In this paper, we use the eigenfunctions of the Laplace-Beltrami operator on the mesh as the basis functions. The cotangent weight scheme [32] without area normalization is employed for the discretization of the Laplace-Beltrami operator, which is less sensitive to volume distortion and results in more compact functional maps. For any real function f on $\mathcal{M}$ represented as $f = \sum_i \tilde{a}_i\phi_i^{\mathcal{M}}$ and $g = T_F(f)$ on $\mathcal{N}$ represented as $g = \sum_i \tilde{b}_j\phi_j^{\mathcal{N}}$, we have the equation:

$$\begin{aligned} g &= \sum_j \tilde{b}_j\phi_j^{N} = T_F(f) = \sum_i \tilde{a}_i T_F(\phi_i^{\mathcal{M}}) \\ &= \sum_i \tilde{a}_i \sum_j c_{ji}\phi_j^{\mathcal{N}} = \sum_j\sum_i \tilde{a}_i c_{ji}\phi_j^{\mathcal{N}}, \end{aligned} \quad (2)$$

which can be rewritten as $b = Ca$ if $a = (\tilde{a}_i)$ and $b = (\tilde{b}_j)$ denote the vectors of coefficients of f and g, respectively. In this way, many constraints of the mapping T become linear in the functional representation, such as descriptor preservation, point or segment correspondences and operator commutativity, and cast enough constraints $b_i = Ca_i$ on the unknown matrix C. According to **Theorem 5.1** in [13], when the underlying map T is isometric, T commutes with the Laplace-Beltrami operator and the corresponding functional matrix C must be orthonormal. Hence the orthogonality and operator commutativity provide additional constraints in this case.

3.2. *Optimization with orthogonality constraints*

It is well known that global intrinsic symmetry is a self-isometric transformation of a shape. It induces an orthonormal functional matrix C commutating with the Laplace-Beltrami operator. As mentioned in Section 3.1, the matrix C can be recovered by casting the following three types of constraints:

$$CA = B, \quad (3)$$

$$CR = RC, \quad (4)$$

$$C^TC = I, \quad (5)$$

where $A = (a_i)$, $B = (b_i)$, and R is the functional matrix induced by the Laplace-Beltrami operator.

In order to find the best transformation in the functional representation satisfying the constraints in *Eq.* 3, 4 and 5, Ovsjanikov et al. [12, 13] and Pokrass et al. [16] estimate an initial functional map by solving a linear system constructed via *Eq.* 3 and 4, in the least squares sense. A post-processing is employed to refine the initial functional map by orthogonalizing it iteratively, in which point-to-point mappings over samples must be established iteratively. The post-processing may also break some existing constraints when refining the initial transformation. Thus, a good initial functional map is important and the computational cost relies on the number of the samples.

In this paper, we employ the optimization method with orthogonality constraints [33] to compute a functional map satisfying all of the constraints directly. The algorithm has lower flops and generates no worse solution than the state-of-the-art methods. Our problem is formulated as follows:

$$min_C \|CA - B\|_F^2 + \lambda\|CR - RC\|_F^2 \ s.t.\ C^TC = I, \quad (6)$$

where I is the identity matrix. We choose the Frobenius norm to ensure the differentiability of the objective function and a non-negative parameter λ to control the influence of operator commutativity. A small λ is used when a shape undergoes some degree of non-isometric deformations. We use $\lambda = 0.1$ for our experiments.

4. Algorithm

Given a nearly self-isometric triangular mesh $\mathcal{M}$, our algorithm takes three stages to establish the underlying dense intrinsic symmetry $T : \mathcal{M} \to \mathcal{M}$. First, we extract a sparse and stable symmetry-invariant point set $\mathcal{V}$ (see the first row of Fig. 3) and establish reliable initial symmetric point pairs $\mathcal{P}_1^S$ from it. The initial regional constraints are constructed from $\mathcal{P}_1^S$. Combining the initial regional constraints with two types of descriptor preservation constraints, denoted as $CA_1 = B_1$, an initial transformation C_1 is computed via *Eq.* 6. Symmetric point pairs over reliable parts of C_1 are established and deemed to be symmetry elector groups $\mathcal{P}_V$. Then, the electors from $\mathcal{P}_V$ cast votes on candidate point pairs outside the reliable parts to establish more symmetric point pairs $\mathcal{P}_2^S$. The final transformation C_2 is solved using *Eq.* 6 with the constraints $CA_1 = B_1$ and $CA_2 = B_2$ constructed from $\mathcal{P}_2^S$. Finally, C_2 is converted to T via a variant of the method described in [13]. The pseudocode of our approach is given in *Algorithm* 1.

Algorithm 1: Properly-constrained Orthonormal Functional Maps for Intrinsic Symmetries

Input: A nearly self-isometric shape $\mathcal{M}$
Output: A point-to-point self-mapping T
1: /*The selection of symmetry electors*/
1.1: $\mathcal{V} \leftarrow ExtractSet(\mathcal{M}, AGD)$;
1.2: $\mathcal{P}_1^S \leftarrow EstablishPairs(\mathcal{V}, HKS, WKS, AGD)$;
1.3: $CA_1 = B_1 \leftarrow BuildConstraints(HKS, WKS, \mathcal{P}_1^S)$;
1.4: $C_1 \leftarrow OptimizeMap(A_1, B_1, R)$;
1.5: $\mathcal{P}_V \leftarrow SelectElectors(C_1, \mathcal{P}_1^S)$.
2: /*The voting scheme for more symmetric point pairs*/
2.1: $\mathcal{P}_2^S \leftarrow ElectorsVoting(\mathcal{P}_V, WKS)$;
2.2: $CA_2 = B_2 \leftarrow BuildConstraints(\mathcal{P}_2^S)$;
2.3: $C_2 \leftarrow OptimizeMap(A_1, B_1, A_2, B_2, R)$.
3: /*The conversion to the dense self-isometry*/
3.1: $T \leftarrow ConvertMap(C_2)$.

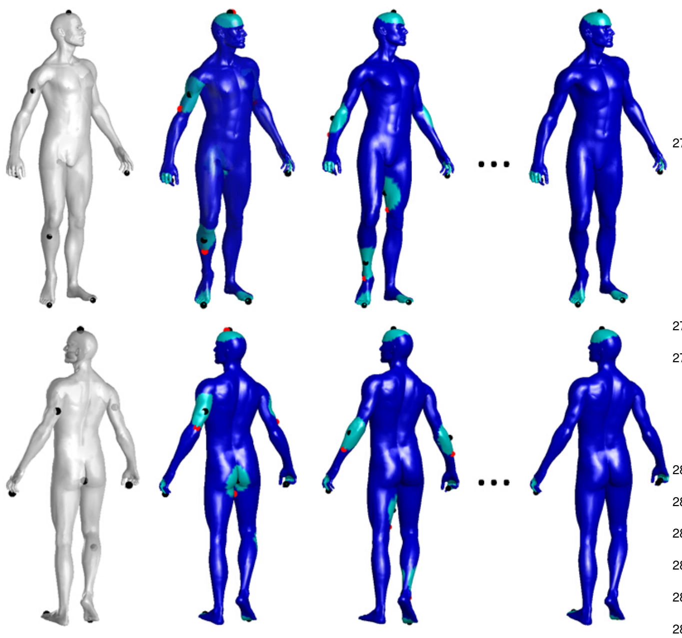

Figure 4: The symmetry-invariant point set $\mathcal{V}$ (the right column) extracted from nine samples (the left column) is shown from two viewpoints (the top row and the bottom row). In each iteration, every sample (the black point) is moved to the maxima (the red point) of AGD within a local region (the cyan region) around it.

4.1. The selection of symmetry electors

4.1.1. The symmetry-invariant point set

A point set $\mathcal{V}$ on a self-isometric shape is a symmetry-invariant set [9] if $\forall\, v \in \mathcal{V}$, $T(v) \in \mathcal{V}$ for all symmetries T. v is a stationary point if $T(v) = v$; (v, v') is a symmetric point pair if $T(v) = v'$. Motivated by the observation that the extremities of the self isometric shape are stable and compose a symmetry-invariant set, we devise a sampling algorithm to extract the extremities based on AGD, of which a larger value indicates the point is closer to the extremities.

Starting from the maxima of AGD, we perform the farthest point sampling algorithm and obtain nine samples. These samples compose a sampling set $\mathcal{V}$ (the left column of Fig. 4), the number of which should be larger than the number of the extremities. The set $\mathcal{V}$ is not necessary to be on the extremities and symmetry-invariant. For each $v_i \in \mathcal{V}$, we move it towards to the extremities iteratively. In each iteration, v_i is moved to the maxima of AGD within a local region around it (the second and third columns of Fig. 4), the radius of which is 0.2 times the maximum geodesic distance. The movement is stopped until v_i is stable or there is a sample $v_j \in \mathcal{V}$ with the AGD value not smaller than the one of v_i. In the latter case, v_i is pruned from $\mathcal{V}$.

4.1.2. The initial symmetric point pairs

We introduce two distortion measures of a mapping over the symmetry-invariant set $\mathcal{V}$. Given a mapping $\mathcal{T} : \mathcal{V} \to \mathcal{V}$, we measure its deviation from isometry by the maximum distortion $d_{iso}(\mathcal{T})$ [34] and average distortion $D_{iso}(\mathcal{T})$ [35] defined as:

$$d_{iso}(\mathcal{T}) = max_{(v_i,v_j)\in\mathcal{T}}\, max_{(v_s,v_t)\in\mathcal{T}'}\, \widetilde{d}_{iso}(v_i, v_j; v_s, v_t), \tag{7}$$

$$D_{iso}(\mathcal{T}) = \frac{1}{|\mathcal{T}|} \sum_{(v_i,v_j)\in\mathcal{T}} \frac{1}{|\mathcal{T}'|} \sum_{(v_s,v_t)\in\mathcal{T}'} \widetilde{d}_{iso}(v_i, v_j; v_s, v_t), \tag{8}$$

where $\mathcal{T}' = \mathcal{T} - (v_i, v_j)$, $\widetilde{d}_{iso}(v_i, v_j; v_s, v_t)$ is the non-isometric distortion between point pairs (v_i, v_j) and (v_s, v_t) defined as:

$$\widetilde{d}_{iso}(v_i, v_j; v_s, v_t) = |d_g(v_i, v_s) - d_g(v_j, v_t)|, \tag{9}$$

where $d_g(\cdot,\cdot)$ is the geodesic metric and is normalized by the maximum geodesic distance over the mesh. A mapping $\mathcal{T}$ is an ambiguous correspondence if the maximum distortion $d_{iso}(\mathcal{T})$ and average distortion $D_{iso}(\mathcal{T})$ are zeros (or approximate to zeros), such as the identity mapping $\mathcal{T}_1$ in Fig. 5 and the flipped mappings $\mathcal{T}_i, i = 2, 3, 4$, in Fig. 5. Once the ambiguous mappings are identified, the symmetry orbit of a point can be extracted directly. An efficient search algorithm is proposed to find the ambiguous mappings over the symmetry-invariant set $\mathcal{V}$, whose stability and sparsity ensure the reliability and feasibility of our search algorithm.

The search algorithm is summarized in the following steps. First, we generate all of the mappings $\{\mathcal{T}\}$ among $\mathcal{V}$ as the search space. Second, according to the local geometric similarity and global distance structure, we prune bad mappings to obtain the ambiguous mappings and the initial symmetric point pairs $\mathcal{P}_1^S$ (Fig. 2 (a)).

A mapping in the search space could be identified as a bad mapping from some perspectives. In our experiment, a mapping is bad if its differences of local descriptors AGD, HKS and WKS are all larger than the corresponding thresholds ϵ_{AGD}, ϵ_{HKS} and ϵ_{WKS}. We do not prune a mapping only relying on one type of descriptors. We also classify a mapping as a bad mapping if its $d_{iso}(\mathcal{T})$ or $D_{iso}(\mathcal{T})$ are larger than prescribed thresholds $\epsilon_{d_{iso}}$ and $\epsilon_{D_{iso}}$.

The above thresholds are determined automatically. Taking the computation of $\epsilon_{D_{iso}}$ as an example (the top row in Fig. 6), we compute $D_{iso}(\mathcal{T})$ for all mappings $\{\mathcal{T}\}$ in the search space and sort the values in ascending order. In this way, we get a parameter curve of $D_{iso}(\mathcal{T})$ and compute its gradient curve. We select the value of $D_{iso}(\mathcal{T})$ corresponding to the first maximum value of the gradient curve (the red point in Fig. 6 (b)) as $\epsilon_{D_{iso}}$. The effectivity of the strategy is attributed to the fact that the distortions of the ambiguous mappings are zeros (or approximate to zeros), which results in a jump between the ambiguous mappings (the black points in Fig. 6 (a)) and the rest mappings. We use the smoothing method in [36] to approximate the prominent structure of the curves. The thresholds ϵ_{AGD}, ϵ_{HKS}, ϵ_{WKS} and $\epsilon_{d_{iso}}$ are determined in the same way.

4.1.3. The initial transformation

For each $(s_i, s_i') \in \mathcal{P}_1^S$, we pick out the geodesic disks centering at s_i and s_i', respectively. The average values of Shape

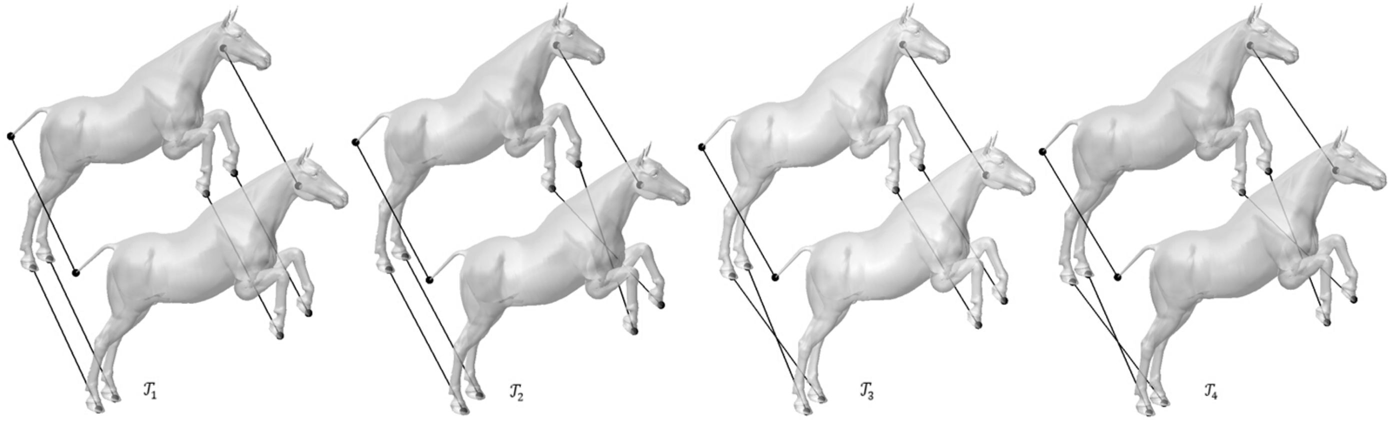

Figure 5: The mappings $\Gamma = \{\mathcal{T}_i | \mathcal{T}_i : \mathcal{V} \to \mathcal{V}, i = 1, 2, 3, 4\}$ could not be distinguished by geometric similarity and distance structure.

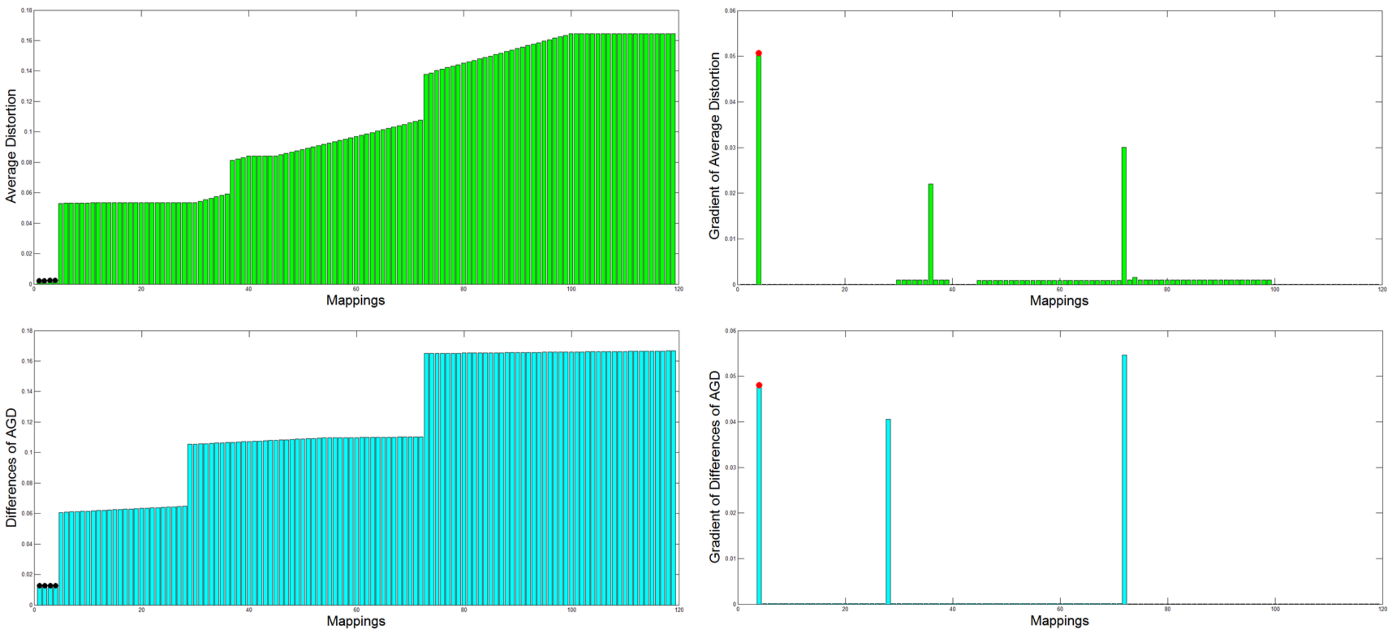

a. The smoothed parameter curves

b. The smoothed gradient curves

Figure 6: Take the horse in Fig. 5 as an example. (a) The smoothed parameter curves of $D_{iso}(\mathcal{T})$ (top) and differences of AGD (bottom) about mappings; (b) The smoothed gradient curves of the smoothed parameter curves. The black points in (a) are the average distortion and differences of AGD of the ambiguous mappings in Fig. 5. The red points in (b) are the first maximum values of the gradient curves.

Diameter Function (SDF) [37] over the geodesic disks are controlled to be less than ϵ_{SDF}, where SDF is normalized by its maximum value. The threshold value ϵ_{SDF} is computed automatically as follows: we segment the model into four clusters according to SDF via k-means. The cluster with the largest SDF value is deemed to be the "body" of the model. We choose the minimal SDF value within this cluster as ϵ_{SDF}. The geodesic disks are divided evenly to obtain geodesic strip pairs as the initial regional constraints. Finally, we project the indicator functions of the geodesic strip pairs onto the orthonormal basis $\{\phi_i^{\mathcal{M}}(x)\}_{i=1}^n$ to obtain linear constraints. Combining the linear constraints with HKS, WKS preservation constraints, we obtain the initial linear constraints $CA_1 = B_1$. The initial transformation C_1 is solved through *Eq.* 6 using $A = A_1$, $B = B_1$. The geodesic disks are deemed to be the reliable parts of C_1.

Ovsjanikov et al. [13] convert the functional map to a point-to-point mapping through searching the image of the delta function centered at each point in all of the delta functions centered at points of $\mathcal{M}$. Different with [13], we limit the search space of points in the reliable parts onto themselves. Moreover, only symmetric points over the reliable parts are generated, denoted as a set of symmetry elector groups $\mathcal{P}_V = \{\mathcal{P}_V^i | i = 1, 2, \cdots, \xi\}$, $\xi = |\mathcal{P}_1^S|$. Each group is a set of symmetric point pairs denoted as symmetry electors $\mathcal{P}_V^i = \{(s_{ij}, s'_{ij}) | j = 1, 2, \cdots\}$, which corresponds to the initial symmetric point pair $(s_i, s'_i) \in \mathcal{P}_1^S$ (as shown in Fig. 2 (b)).

4.2. The voting scheme for more linear constraints

To make our method more robust to non-isometric deformations, we propose a voting scheme to construct more constraints on parts far away from the shape extremities, such as the torso. First, we initialize candidate point pairs from the local maxima of Wave Kernel Signature outside the reliable parts. Second, for each candidate point pair (v_s, v_t), we select the nearest symmetry elector group $\mathcal{P}_V^i$, which is measured by the average geodesic distance between (v_s, v_t) and the initial symmetric point pair (s_i, s_i'). Then each elector $(s_{ij}, s_{ij}') \in \mathcal{P}_V^i$ casts a vote on (v_s, v_t) if the degree of asymmetry of (v_s, v_t) is less than a prescribed threshold δ_1, which is 0.07 times the maximum geodesic distance in our implementation. The asymmetry of point pair (v_s, v_t) evaluated by the elector (s_{ij}, s_{ij}') is defined as follows:

$$asym((v_s, v_t), (s_{ij}, s_{ij}')) = \\ max(|d_g(v_s, s_{ij}) - d_g(v_t, s_{ij}')|, |d_g(v_s, s_{ij}'), d_g(v_t, s_{ij})|). \quad (10)$$

After the voting procedure, we prune the candidate point pairs whose votes are less than a half of the number of their corresponding symmetry elector group, and filter bad point pairs according to the local geometric similarity of WKS, which is 0.15. Finally, we obtain more symmetry point pairs $\mathcal{P}_2^S$ (Fig. 2 (c)). More regional constraints are obtained from $\mathcal{P}_2^S$ through choosing a small geodesic disk per point, and converted to linear constraints $CA_2 = B_2$. Combining $CA_1 = B_1$ with $CA_2 = B_2$, the final transformation C_2 is solved using *Eq.* 6. We convert C_2 to the dense intrinsic symmetry T using a limited search space, instead of the whole space of delta functions centered at points of $\mathcal{M}$ employed in [13]. In practice, we search the symmetric points of parts within the symmetric parts, and search their symmetric points of torso within the torso itself, which results in accurate symmetry map and reduces the computational cost simultaneously.

5. Experimental results

In this section, we evaluate our method on two datasets of the intrinsic symmetry benchmark [9]. The TOSCA dataset [38] has 80 shapes with approximate intrinsic symmetries in 9 classes. The SCAPE dataset [39] contains 71 shapes in 1 class. Similar to [12], we improve the quality of the benchmark, and increase the number of it to 200 uniformly distributed points for each shape class in TOSCA, denoted as an augmented benchmark $\mathcal{U} = \{u_i | i = 1, 2, \cdots, 200\}$. $\mathcal{U}$ is computed via the global extrinsic reflection symmetry of the undeformed shape in each class. The shape class "gorilla" is excluded because there is no undeformed version of it. For each point $u_i \in \mathcal{U}$, we denote the geodesic distance between $T(u_i)$ and u_i', which is the ground-truth correspondence of u_i, as the geodesic error. The two evaluation metrics in [9] are employed:

- **Correspondence rate:** The percentage of points in $\mathcal{U}$ for which the geodesic error is less than a distance threshold ε.
- **Mesh rate:** The percentage of shapes for which the correspondence rate is above a threshold β.

	Corr Rate (%)			Mesh Rate		
TOSCA	LS	IR	Our	LS	IR	Our
Cat	61.8	77.6	**86.1**	2/11	6/11	**10/11**
Centaur	74.0	82.1	**94.6**	3/6	3/6	**6/6**
David	65.9	69.2	**87.8**	3/7	3/7	**7/7**
Dog	82.0	86.6	**91.7**	7/9	7/9	**8/9**
Horse	88.2	92.8	**95.1**	7/8	7/8	7/8
Michael	71.2	73.1	**90.7**	12/20	13/20	**17/20**
Victoria	74.3	80.6	**95.4**	5/11	8/11	**11/11**
Wolf	98.8	99.8	**100**	3/3	3/3	3/3
Average	74.2	80.0	**91.7**	42/75	50/75	**69/75**

Table 1: Average Correspondence Rate and Mesh Rate of the results solved via LS, IR [13] and our optimization 6. The evaluations are conducted on $\mathcal{U}$ for all classes of TOSCA, except for "gorilla". The transformations solved by LS are the initializations of IR and the rates are listed in the left column. The rates of IR and our optimization method are listed in the middle and the right columns.

In our experiments, $\varepsilon = \sqrt{\frac{area(M)}{20\pi}}$, $\beta = 75\%$, which are the same as [9].

5.1. Comparison of our optimization with the iterative refinement method

In order to evaluate the effectivity of our optimization with orthogonality constraints, we compare the functional maps generated by *Eq.* 6 with the ones obtained by the linear system (LS) and the iterative refinement method (IR) [13]. Ovsjanikov et al. [13] estimate an initial functional map from LS constructed via *Eq.* 3 and 4, and refine it using IR. To be fair, we use the same constraints in *Eq.* 3 and 4. The conversion procedure from functional maps to point-to-point mappings over the augmented benchmark $\mathcal{U}$ is the same, too. We solve *Eq.* 6 with $\lambda = 1$ to factor out the affection of it and run IR for 20 iterations as mentioned in [13]. As the statistics in Tab. 1 illustrate, our optimization with orthogonality constraints gives much better functional maps than IR, since we search the best transformation satisfying all of the constraints directly. It works well even when the shapes undergo some non-isometric deformations (see the last two rows of Fig. 7). Although IR improves the initial estimation from LS in general, it may break good correspondences in some degree during removing the bad ones, as shown in the top row of Fig. 7. We also find that LS may not provide a good initial estimation for the shapes with moderate non-isometric deformations. Without the good initial estimation, IR fails to perform well, as illustrated in the last two rows of Fig. 7.

5.2. Comparison of our method with and without the voting scheme

In practice, many natural objects or man-made models are not perfectly symmetric and often undergo some degree of non-isometric deformations. Thus inadequate regional constraints, such as initial regional constraints, with the two types of feature preservation constraints (HKS, WKS) are not enough to generate robust self-isometry using the function map framework (Tab. 2). As the last two rows of Fig. 8 show, more regional

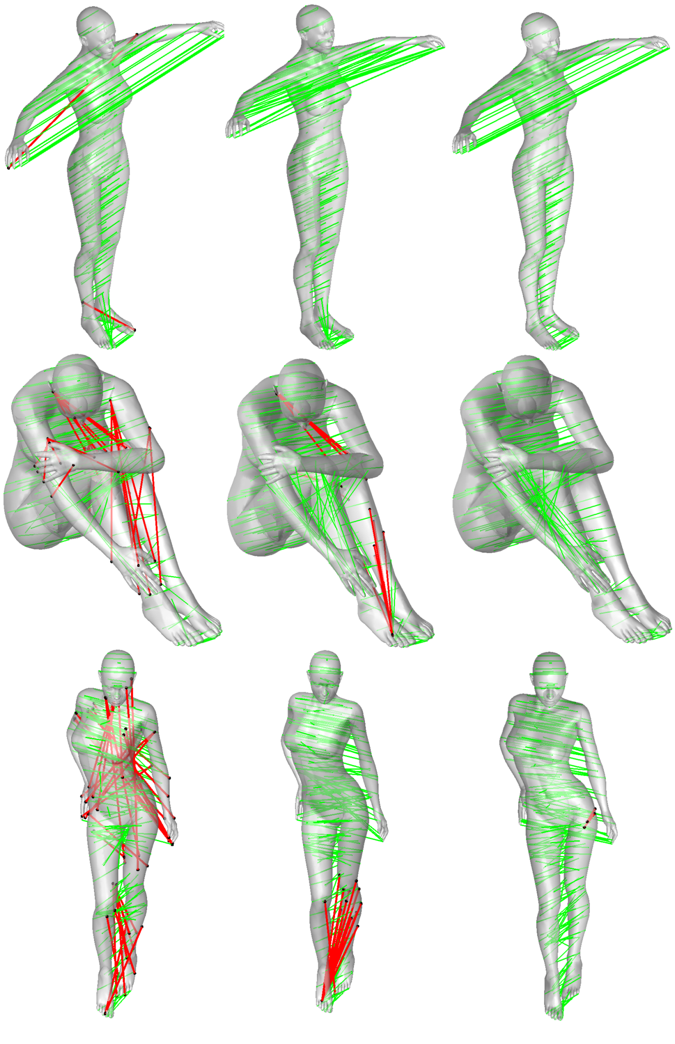

Figure 7: Global intrinsic symmetries of LS, IR [13] and our method on $\mathcal{U}$ are listed from left to right. Green lines link point pairs with small geodesic errors while red lines link pairs with larger geodesic errors. The average distortions of meshes from the top row to the bottom row are 0.0000, 0.0090 and 0.0121.

constraints detected by the voting process contribute to generate reasonable results when the centaurs undergo some non-isometric deformations.

5.3. Comparison with the state-of-the-art methods

To evaluate the quality of our symmetry detection algorithm, we compare the results on TOSCA and SCAPE datasets with the state-of-the-art methods, Möbius transformations (MT) [9] and Blended Intrinsic Maps (BIM) [30], in Tab. 3 and Tab. 4, respectively. Some visual comparisons with them are illustrated in Fig. 9. The comparisons are based on the manually selected ground-truth set from MT instead of the aforementioned augmented benchmark $\mathcal{U}$. The average correspondence rates and mesh rates of our method are 95.1% and 76/79 on TOSCA (Tab. 3), and 91.7% and 69/71 on SCAPE (Tab. 4), which are significant improvements over MT. Compared to BIM, we improve the statistics on SCAPE dataset because BIM may map the front of a human to the back (the middle column in Fig. 9). The quality of SCAPE meshes is not as good as the one of

	Corr Rate (%)		Mesh Rate	
TOSCA	Without	With	Without	With
Cat	76.2	**86.5**	6/11	**10/11**
Centaur	80.5	**94.6**	5/6	**6/6**
David	84.1	**88.1**	5/7	**7/7**
Dog	84.8	**91.8**	7/9	**8/9**
Horse	82.8	**95.0**	6/8	**7/8**
Michael	77.0	**90.7**	10/20	**17/20**
Victoria	84.3	**95.5**	9/11	**11/11**
Wolf	97.5	**100**	3/3	3/3
Average	81.3	**91.8**	51/75	**69/75**

Table 2: Average Correspondence Rate and Mesh Rate of our method without the voting scheme and our full method. The evaluations are conducted on $\mathcal{U}$ for all shape classes of TOSCA, except for ”gorilla”. The rates of our full method are listed after the method without the voting scheme.

	Corr Rate(%)			Mesh Rate		
TOSCA	MT	BIM	Our	MT	BIM	Our
Cat	66	**93.7**	90.9	6/11	10/11	10/11
Centaur	92	**100**	96.0	6/6	6/6	6/6
David	82	**97.4**	94.8	4/7	7/7	7/7
Dog	91	**100**	93.2	8/9	**9/9**	8/9
Horse	92	**97.1**	95.2	8/8	**8/8**	7/8
Michael	87	**98.9**	94.6	15/20	20/20	20/20
Victoria	83	**98.3**	98.7	7/11	11/11	11/11
Wolf	100	100	100	3/3	3/3	3/3
gorilla	–	98.9	98.9	–	4/4	4/4
Average:	85	**98.02**	95.1	57/75	**78/79**	76/79

Table 3: Comparison of MT [9], BIM [30] and our method on TOSCA. The statistics are based on the manually selected ground-truth set used in [9].

	Corr Rate(%)			Mesh Rate		
SCAPE	MT	BIM	Our	MT	BIM	Our
Average:	82	84.8	**91.7**	51/71	54/71	**69/71**

Table 4: Comparison of MT [9], BIM [30] and our method on SCAPE. The statistics are based on the manually selected ground-truth set used in [9].

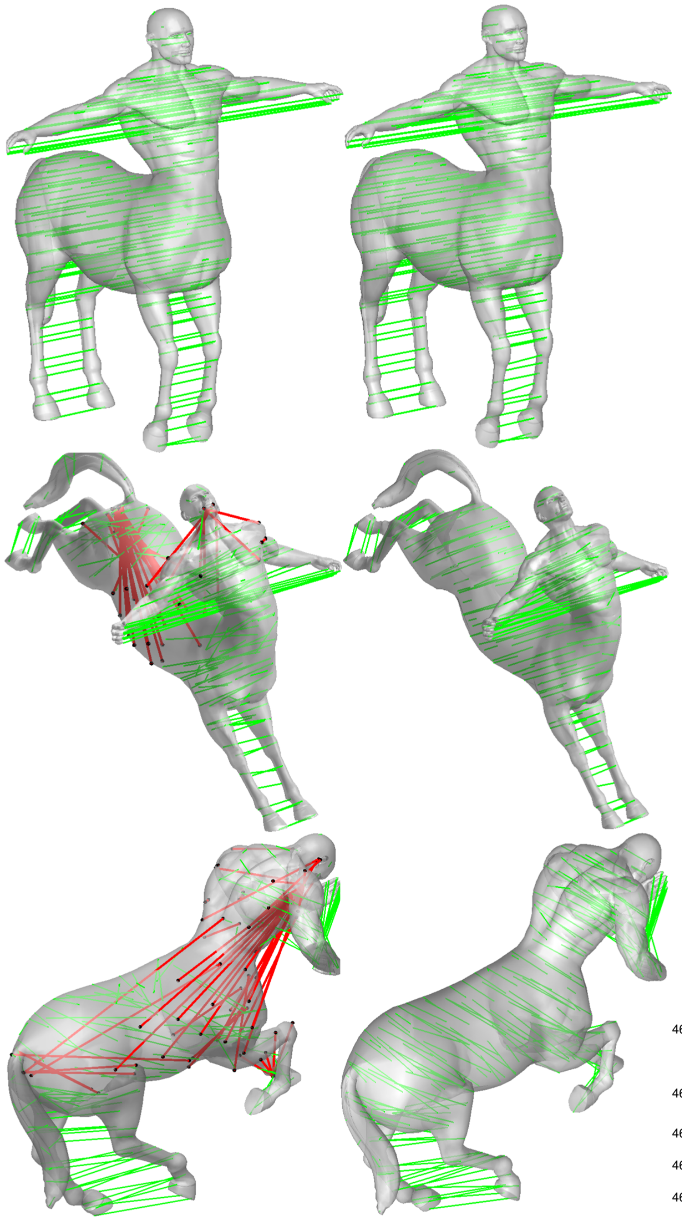

Figure 8: Comparison of our method without the voting scheme (the left column) and our full method (the right column) on $\mathcal{U}$. The average distortions of the ground-truth correspondences of the shapes from the top row to the bottom row are 0.0000, 0.0076 and 0.0140, respectively. Green lines link point pairs with small geodesic errors while red lines for the other.

TOSCA meshes and this may explain the decrease in performance for our algorithm and BIM on SCAPE. The performance of BIM drops more for the mid-edge uniformization technique employed. BIM provides better statistics on TOSCA than our algorithm, but it suffers from the running time issue. Our Matlab implementation takes 67.1 minutes to compute the global symmetries for all TOSCA meshes, while the BIM's C++ implementation takes 365.5 minutes. More results of our method on TOSCA and SCAPE datasets are presented on $\mathcal{U}$ in Fig. 12. Our algorithm generalizes well to other classes of shapes with extremities in addition to humans and animals. We depict some results on models from the SHREC 2007 Watertight Benchmark [40] in Fig. 10.

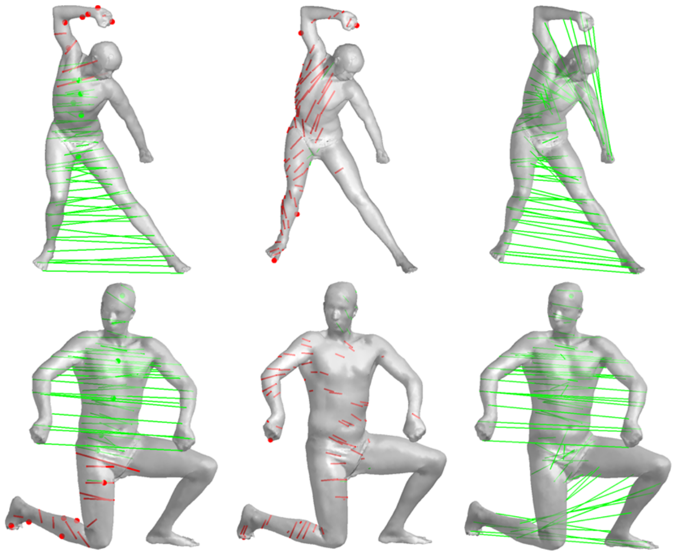

Figure 9: Comparison of the MT [9] (the left column), BIM [30] (the middle column) and our method (the right column). Points indicate they are mapped to themselves. Green lines and points correspond to point pairs with small geodesic errors while red lines and points for the other.

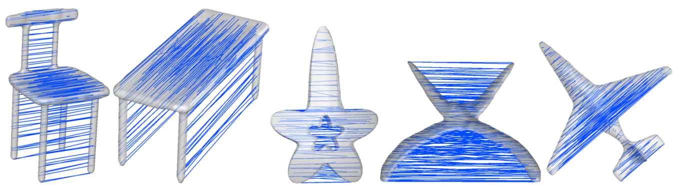

Figure 10: Global intrinsic symmetries of other shapes with extremities.

5.4. *Limitations*

While our method handles shapes with moderate non-isometric distortion, it still has some limitations. The first limitation is that our method could not detect the intrinsic symmetries of models without extremities, or without symmetric point pairs in the symmetry-invariant point set. As shown in Fig. 11 (a), our method fails because the symmetry-invariant point set of the vase only contains stationary points (the black points), and provides no initial symmetric point pairs. Moreover, if the reliable parts for the subsequent voting scheme are insufficient, we may get unsatisfactory results (Fig. 11 (b) and (c)). Hence we will find a more general scheme to construct sufficient initial constraints. The second limitation is that we handle only reflectional symmetry in this paper. We plan to extend our method to explore more general cases of symmetries in the future.

6. Conclusion

In this paper, we introduce a novel intrinsic symmetry detection method. Instead of propagating the sparse correspondence to the entire shape using geodesic distance, the compact functional map framework is leveraged. We design an initialization procedure to extract sparse and reliable symmetric point pairs

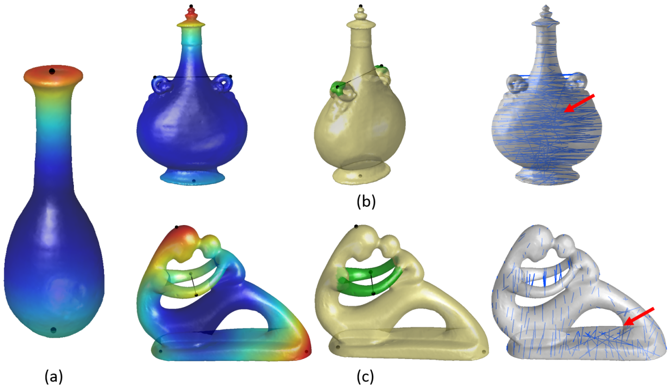

Figure 11: The failures of our algorithm. (a) The symmetry-invariant point set of the vase only contains stationary points (the black points). (b) and (c) The reliable regions (the green regions) of initial symmetric pairs (the black lines) are insufficient for the subsequent voting scheme, which results in unsatisfactory results (red arrow).

from the extremities of the model, and a voting procedure to extract more symmetric point pairs distributed over the entire shape. The symmetric point pairs are then employed to construct a set of regional constraints. Finally, we formulate the problem as an optimization with descriptor, regional and orthogonality constraints simultaneously. The functional representation, efficient optimization method and effective regional constraints together make our method a faster, automatic and robust implementation. Experimental results on the symmetry detection benchmark exhibit the improved accuracy of our method for a large variety of object types with moderate deviations from perfect symmetry.

Acknowledgement

The authors sincerely thank all the anonymous reviewers for their valuable suggestions. Thanks to AIM@SHAPE for the SHREC 2007 Watertight Models. Xiuping Liu is supported by the NSFC Fund (Nos. 61173102, 61370143). Risheng Liu is supported by the NSFC Fund (No. 61300086), the Fundamental Research Funds for the Central Universities (No. DUT12RC(3)67) and the Open Project Program of the State Key Laboratory of CAD&CG, Zhejiang University, China (No. A1404). Hui Wang is supported by the Natural Science Foundation of Hebei Province (F2014210127) and Scientific Research Foundation of Shijiazhuang Tiedao University (Z110205). Junjie Cao is supported by the NSFC Fund (No. 61363048) and the Fundamental Research Funds for the Central Universities (Nos. DUT13RC206, DUT13JS04).

References

[1] M. Kazhdan, T. Funkhouser, S. Rusinkiewicz, Symmetry descriptors and 3d shape matching, in: Symposium on Geometry Processing, 2004, pp. 115–123.

[2] P. Simari, E. Kalogerakis, K. Singh, Folding meshes: Hierarchical mesh segmentation based on planar symmetry, in: Symposium on Geometry Processing, 2006, pp. 111–119.

[3] S. Thrun, B. Wegbreit, Shape from symmetry, in: IEEE International Conference on Computer Vision, Vol. 2, 2005, pp. 1824–1831.

[4] J. Podolak, A. Golovinskiy, S. Rusinkiewicz, Symmetry-enhanced remeshing of surfaces, in: Symposium on Geometry Processing, 2007, pp. 235–242.

[5] N. J. Mitra, M. Pauly, M. Wand, D. Ceylan, Symmetry in 3d geometry: Extraction and applications, Computer Graphics Forum 32 (6) (2013) 1–23.

[6] N. J. Mitra, L. J. Guibas, M. Pauly, Partial and approximate symmetry detection for 3d geometry, ACM Transactions on Graphics 25 (3) (2006) 560–568.

[7] J. Podolak, P. Shilane, A. Golovinskiy, S. Rusinkiewicz, T. Funkhouser, A planar-reflective symmetry transform for 3d shapes, ACM Transactions on Graphics 25 (3) (2006) 549–559.

[8] D. Raviv, A. M. Bronstein, M. M. Bronstein, R. Kimmel, Symmetries of non-rigid shapes, in: IEEE International Conference on Computer Vision, 2007, pp. 1–7.

[9] V. G. Kim, Y. Lipman, X. Chen, T. Funkhouser, Möbius transformations for global intrinsic symmetry analysis, Computer Graphics Forum 29 (5) (2010) 1689–1700.

[10] Y. Lipman, T. Funkhouser, Möbius voting for surface correspondence, ACM Transactions on Graphics 28 (3) (2009) 72.

[11] M. Hilaga, Y. Shinagawa, T. Kohmura, T. L. Kunii, Topology matching for fully automatic similarity estimation of 3d shapes, in: Proceedings of SIGGRAPH, 2001, pp. 203–212.

[12] M. Ovsjanikov, Q. Mérigot, V. Pătrăucean, L. Guibas, Shape matching via quotient spaces, Computer Graphics Forum 32 (5) (2013) 1–11.

[13] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, L. Guibas, Functional maps: A flexible representation of maps between shapes, ACM Transactions on Graphics 31 (4) (2012) 30.

[14] J. Sun, M. Ovsjanikov, L. Guibas, A concise and provably informative multi-scale signature based on heat diffusion, Computer Graphics Forum 28 (5) (2009) 1383–1392.

[15] M. Aubry, U. Schlickewei, D. Cremers, The wave kernel signature: A quantum mechanical approach to shape analysis, in: IEEE International Conference on Computer Vision Workshops, 2011, pp. 1626–1633.

[16] J. Pokrass, A. M. Bronstein, M. M. Bronstein, P. Sprechmann, G. Sapiro, Sparse modeling of intrinsic correspondences, Computer Graphics Forum 32 (2pt4) (2013) 459–468.

[17] R. Gal, D. Cohen-Or, Salient geometric features for partial shape matching and similarity, ACM Transactions on Graphics 25 (1) (2006) 130–150.

[18] A. Martinet, C. Soler, N. Holzschuch, F. X. Sillion, Accurate detection of symmetries in 3d shapes, ACM Transactions on Graphics 25 (2) (2006) 439–464.

[19] R. M. Rustamov, Augmented symmetry transforms, in: IEEE International Conference on Shape Modeling and Applications, 2007, pp. 13–20.

[20] D. Raviv, A. M. Bronstein, M. M. Bronstein, R. Kimmel, G. Sapiro, Diffusion symmetries of non-rigid shapes, in: Proceedings of International Symposium on 3D Data Processing, Visualization and Transmission, 2010.

[21] M. Ovsjanikov, J. Sun, L. Guibas, Global intrinsic symmetries of shapes, Computer Graphics Forum 27 (5) (2008) 1341–1348.

[22] K. Xu, H. Zhang, A. Tagliasacchi, L. Liu, G. Li, M. Meng, Y. Xiong, Partial intrinsic reflectional symmetry of 3d shapes, ACM Transactions on Graphics 28 (5) (2009) 138:1–138:10.

[23] S. Wang, Y. Wang, M. Jin, X. D. Gu, D. Samaras, Conformal geometry and its applications on 3d shape matching, recognition, and stitching, IEEE Transactions on Pattern Analysis and Machine Intelligence 29 (7) (2007) 1209–1220.

[24] A. Berner, M. Bokeloh, M. Wand, A. Schilling, H.-P. Seidel, A graph-based approach to symmetry detection, in: Proceedings of the IEEE VGTC Workshop on Volume Graphics 2008, 2008, pp. 1–8.

[25] M. Bokeloh, A. Berner, M. Wand, H.-P. Seidel, A. Schilling, Symmetry detection using feature lines, Computer Graphics Forum 28 (2) (2009) 697–706.

[26] R. Lasowski, A. Tevs, H.-P. Seidel, M. Wand, A probabilistic framework for partial intrinsic symmetries in geometric data, in: IEEE International Conference on Computer Vision, 2009, pp. 963–970.

[27] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. Van Kaick, A. Tagliasacchi, Deformation-driven shape correspondence, Computer Graphics Forum 27 (5) (2008) 1431–1439.

[28] O. K.-C. Au, C.-L. Tai, D. Cohen-Or, Y. Zheng, H. Fu, Electors voting for fast automatic shape correspondence, Comput. Graph. Forum 29 (2)

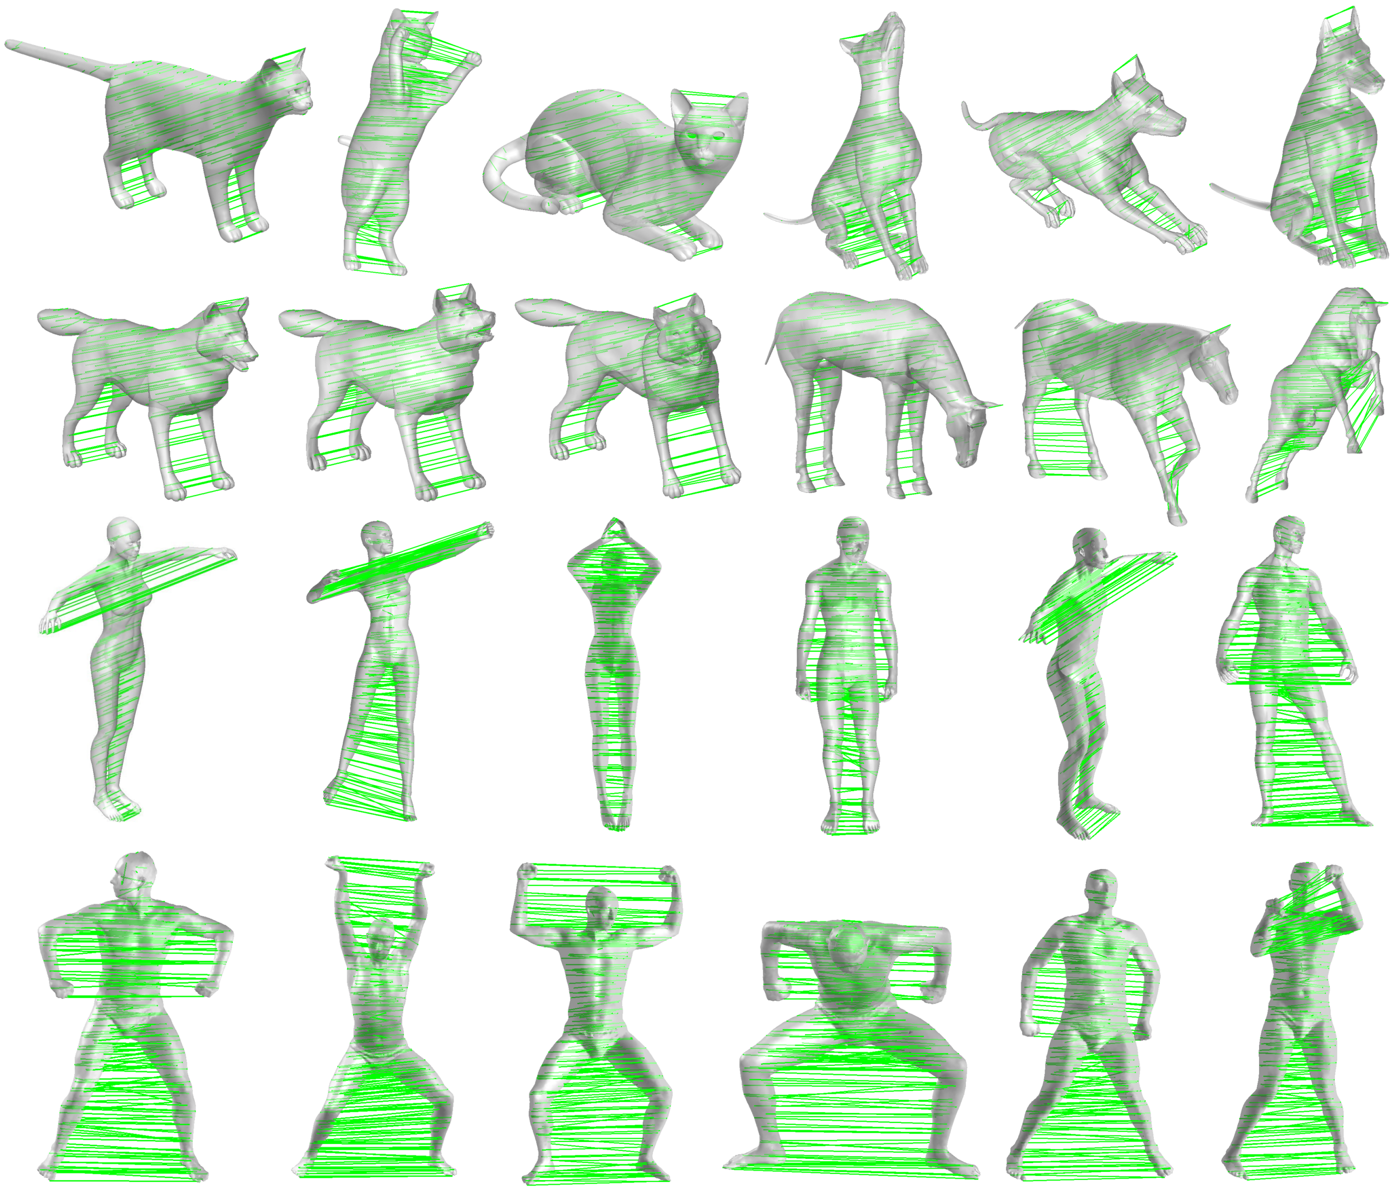

Figure 12: More results of our method on the augmented benchmark $\mathcal{U}$. The models in the first three rows are from the TOSCA dataset while the ones in the last row are from the SCAPE dataset.

(2010) 645–654.

[29] R. M. Rustamov, Laplace-beltrami eigenfunctions for deformation invariant shape representation, in: Symposium on Geometry Processing, 2007, pp. 225–233.

[30] V. G. Kim, Y. Lipman, T. Funkhouser, Blended intrinsic maps, ACM Transactions on Graphics 30 (4) (2011) 79:1–79:12.

[31] T. Liu, V. G. Kim, T. Funkhouser, Finding surface correspondences using symmetry axis curves, Computer Graphics Forum 31 (5) (2012) 1607–1616.

[32] M. Meyer, M. Desbrun, P. Schröder, A. H. Barr, Discrete differential-geometry operators for triangulated 2-manifolds, in: Visualization and mathematics III, 2003, pp. 35–57.

[33] Z. Wen, W. Yin, A feasible method for optimization with orthogonality constraints, Mathematical Programming 142 (1-2) (2013) 397–434.

[34] A. M. Bronstein, M. M. Bronstein, R. Kimmel, Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching, Proceedings of the National Academy of Sciences of the United States of America 103 (5) (2006) 1168–1172.

[35] Y. Sahillioglu, Y. Yemez, Minimum-distortion isometric shape correspondence using em algorithm, IEEE Transactions on Pattern Analysis and Machine Intelligence 34 (11) (2012) 2203–2215.

[36] L. Xu, C. Lu, Y. Xu, J. Jia, Image smoothing via l0 gradient minimization, ACM Transactions on Graphics 30 (6) (2011) 174:1–174:12.

[37] L. Shapira, S. Shalom, A. Shamir, D. Cohen-Or, H. Zhang, Contextual part analogies in 3d objects, International Journal of Computer Vision 89 (2-3) (2010) 309–326.

[38] A. M. Bronstein, M. M. Bronstein, R. Kimmel, Numerical geometry of non-rigid shapes, Springer Publishing Company, Incorporated, 2008.

[39] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, J. Davis, Scape: shape completion and animation of people, ACM Transactions on Graphics 24 (3) (2005) 408–416.

[40] D. Giorgi, S. Biasotti, L. Paraboschi, Shrec: Shape retrieval contest: Watertight models track (2007).