

Feature Correspondences using Morse Smale Complex

Wei Feng · Jin Huang · Tao Ju · Hujun Bao

Received: date / Accepted: date

Abstract Establishing corresponding features on two non-rigidly deformed 3D surfaces is a challenging and well-studied problem in computer graphics. Unlike previous approaches that constrain the matching between feature pairs using isometry-invariant distance metrics, we constrain the matching using a discrete connectivity graph derived from the Morse-Smale Complex of the Auto Diffusion Function. We observed that the graph remains stable even for surfaces differing by topology or by significant deformation. This algorithm is simple to implement and efficient to run. When tested on a range of examples, our algorithm produces comparable results with state-of-art methods on surfaces with strong isometry but with greatly improved efficiency, and often gets better correspondences on surfaces with larger shape variances.

Keywords Point matching · Correspondence · Morse-Smale complex

1 Introduction

Identifying corresponding feature points on two surfaces is one of the fundamental problems in computer graphics and vision, and has important applications such as registering anatomical structures in medical imaging. While humans are adept at this task, it is extremely challenging for computers to achieve the same accuracy

and efficiency, particularly for shapes that differ by a non-rigid deformation.

There have been significant efforts in the past in developing automated algorithms for non-rigid correspondence finding (see reviews in the next section). These algorithms have been mostly successful for surfaces that are either close-by in Euclidean space, or differ by an *isometric* deformation (e.g., articulation) that preserves geodesic distances between surface points. However, efficiently and accurately finding correspondences for surfaces that exhibit notable differences in shape (e.g., the dinosaurs in Figure 1 top) and even topology (e.g., the hands in Figure 1 bottom) is still challenging.

In this work, we present a new algorithm for finding feature correspondences for non-rigidly deformed surfaces, which we have observed to perform robustly and efficiently in various tests (e.g., Figure 1). Like several previous methods [1,13,28,27,6], we constrain the matching problem using spatial relations among feature points. However instead of using isometry-invariant distances (e.g., geodesic distances), we use a topological graph that encodes the connectivity among feature points as our constraints. The graph, as well as the feature points, are derived from the Morse-Smale Complexes (MSC) of an isometry-invariant function [11,29] over the surface. We have observed that the MSC of such function is stable even for surfaces with fairly large shape variances, where typical isometry-invariant distances are not sufficient to constrain the matching. Also, the MSC-based graph contains a sparse (but essential) set of pair-wise connectivity among the feature points, which has better discriminative power and enables more efficient computation.

Our main contribution is that we propose a new connectivity between feature points based on MSC, which is more discriminative and stable under significant de-

W. Feng · J. Huang · H.J. Bao
CAD&CG lab, Zhejiang University, China
E-mail: hj@cad.zju.edu.cn

T. Ju
Department of Computer Science and Engineering
Washington University in St. Louis



Fig. 1 Feature correspondences on surfaces with varying shape (top) and topology (bottom).

formation. We also propose a local correction process that reduces errors due to symmetry, which is common in similar approaches. This approach is efficient, simple to implement, and has no restriction on the shape or topology of the surface.

We test the algorithm on various examples including those in the SHREC benchmarks [12,3], and compare it with several state-of-art methods. We have observed that our algorithm gives comparable results as these methods on models exhibiting strong isometry, and often outperforms on models with large shape variances. In addition, our algorithm completes in a matter of seconds, whereas other methods would typically take minutes to compute.

2 Related Work

We are finding point correspondences between surfaces which we focus here. More types of correspondence problem please refer to a recent survey ([15]).

Feature point identification There are many methods for identifying salient point locations on surfaces for the purpose of matching. To withstand deformations such as articulated motion, typical choice of features

are critical points of some isometry-invariant functions over the surface, such as eigenfunctions of the Laplace-Beltrami operator [25], the average geodesic distance function [16,17], and the Auto Diffusion Function (ADF)[11] (a.k.a the Heat Kernel Signature (HKS)[29]). In this paper, we consider the critical points of the ADF (as done similarly in several recent works like [27,21,28]), due to its observed stability on surfaces with similar shapes. However, other than using only maxima, we use all critical points and we have observed that some of them are still stable under significant deformation.

Feature correspondences Finding correspondences between two sets of feature points is an ill-posed problem, since the local shape at different feature points can often be similar. As a result, additional constraints on the global relation between the features must be imposed to make the problem trackable.

Under the typical assumption of isometric deformation (e.g., articulated motion), different kinds of constraints have been used in previous works. Zhang et al. [33] evaluates the quality of correspondence by the distortion of the shape after a deformation in the Euclidean space guided by the corresponding points. In a similar vein, Lipman and Funkhouser [19] and later Zeng et al. [32] consider deviation of the features after a Mobius transformation in the parametrization space defined by three feature points.

Another common type of isometry-invariant constraint is geodesic distance. That is, a correspondence that best maintains the pair-wise geodesic distances among feature points is preferred. Methods exist that either embed the shapes in a higher-dimensional space that captures the geodesic distances [14], or impose the geodesic constraints during matching [1,13,28,30]. Note that Huang et al. [13] uses a similar spectral method [18] in their graph-matching framework, although a different graph is adopted in our work.

Since geodesic distances can be unstable under topological changes of the surface, they are not suitable for matching features on similar-looking surfaces but with different genus. To address this deficiency, *diffusion distances* constraints have been recently adopted in place of geodesic distances for more robust feature matching [27,6]. Diffusion distances, first introduced by Coifman and Lafon [8], describes the probability of a path from one point to another on the surface via random walks. Such distances are observed to be less sensitive to topological changes than geodesic distances.

Rather than using isometry-invariant distance metrics as in previous works, we explore the use of a topological connectivity among the feature points as constraints for matching. We observe that such connecti-

ty, which is invariant under isometry, is also quite stable even when geodesic or diffusion distances are not, and is more discriminative.

MSC and geometry processing Morse Smale Complex (MSC) comes from Morse theory ([20]) which characterizes full behavior of scalar field defined on manifold, especially the topological structure of the scalar field. It has been found uses in many surface processing tasks, such as quadrangulation [9], scalar field smoothing [31], and segmentation [24]. To the best of our knowledge, MSC has not yet been used to for finding feature correspondences on surfaces.

3 Overview

Our approach is to search coarse correspondences between two surfaces possibly differing by significant deformation. The corresponding points capture the features of the surfaces, thus they are meaningful for some applications. For example, they can be used as markers for establishing dense correspondences or for cross-parametrization. This approach is efficient and is not limited to surfaces of particular genus or type.




Fig. 2 An incomplete graph versus a complete graph.

Our key idea is to propose discriminative and stable connectivity between points, which are derived from Morse-Smale Complex (MSC), to drive the matching of these points as a graph matching problem. There are two major advantages for this connectivity: firstly, it captures the essential relations between the points, which is more discriminative, compared to the other two commonly used connectivity (geodesic or diffusion distance), which usually generate a *complete* graph. As an example, there will be many matching results for a complete graph to match to another complete graph of the same number of nodes as in Figure 2 right, regardless of the differences of the points. For example, we can rotate the graph by 60° and have another matching result. On the other hand, our connectivity is sparse which generates an *incomplete* graph, e.g. as in Figure 2

left, and the result is unique if it is matched to a graph of the same connectivity, if only rotation is allowed. The main improvements are from the fact that differences of the connectivity (e.g. the valence of the nodes) would distinguish correct candidates from incorrect ones, or at least reduce the candidate numbers. As a real example, we are able to match a cat and a dog (Figure 3) by setting all node differences ineffective ($H(a) = 1.0$, Section 5), in which case we use only the connectivity and node type, while using geodesic or diffusion distance with node type we could not get the correct result. Note that though some other results could be matched with this configuration, we only use it here. Secondly, this connectivity is less sensitive to non-isometric deformations, which allows us to match surfaces differing by significant deformation, providing that the distribution of the points remains. On the contrary, the geodesic or diffusion distance degenerate much faster to non-isometric deformations. The comparisons are given in the result section.

A key issue of this approach is to sample the points *at* the features before constructing the connectivity. Fortunately, we are using Auto Diffusion Function (ADF) to construct the MSC, whose critical points are claimed to capture the features ([11]). More details and the way we construct the connectivity (thus the graph) will be given in the next section. The matching algorithm is given in Section 5, followed by comprehensive evaluations and comparisons in Section 6.




Fig. 3 Matching a cat and a dog with connectivity.

4 Feature Graph

Our key idea is to constrain the matching of feature points with a connectivity graph that is stable for surfaces with similar feature distribution. However this approach relies heavily on how to sample the points, which is highly related to an isometry invariant scalar function (ADF). We will briefly review ADF and how we sample points first. Then we detail the construction of

the *feature graph*, which includes the construction of the graph and the attributes assignment for the graph elements.

4.1 Auto Diffusion Function and Point Sampling

The Auto Diffusion Function (ADF) was independently proposed by Gebal et al. [11] and Sun et al. [29]. It is a solution to the heat equation

$$(\Delta_x + \partial_t)u(x, t) = 0$$

where Δ_x denotes the Laplace-Beltrami operator (LBO) acting on spatial variable x and $t \in [0, \infty)$ is the time variable. The initial condition for this equation is $u(x, 0) = \delta(x)$ where $\delta(x)$ is the Dirac delta function at position x . This solution K can be expressed in the terms of LBO eigensolutions as:

$$K(x, x, t) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i^2(x)$$

where $0 = \lambda_0 \leq \lambda_1 \leq \dots$ are the eigenvalues of LBO and ϕ_i are the corresponding eigenfunctions. To make this function scale invariant, it is normalized by dividing the exponential by the second eigenvalue and ADF is defined as:

$$ADF_t(x) = K(x, x, \frac{t}{\lambda_1}) = \sum_{i=0}^{\infty} e^{-t \frac{\lambda_i}{\lambda_1}} \phi_i^2(x) \quad (1)$$

This normalization is valid as long as the object is complete and for partial objects a more sophisticated normalization method can be adapted ([7]).

In a nutshell, ADF computes how much heat is left at a surface point after diffusing a unit amount of heat at this point to the rest of the surface over a given period of time t . The local extrema of ADF are typically located at tips of features on the surface, and the scale of these features corresponds well to the value of t ([11]). In our experiments, we typically use $t = 0.5$ for coarse scale features and eigens with $e^{-t \frac{\lambda_i}{\lambda_1}} < 0.01$ are used ([11]). One could use as small t as possible to get more critical points, but at the risk of getting poor correspondences or damaging the connectivity we propose. As we show different choices of t in Figure 5, points near eyes of the two heads of $t = 0.2$ (second row) are poorly located because for this finer scale they correspond to different features, and points at the cheek (in rectangle) miss corresponding points which will create different connectivity for the two shapes and make the connectivity less discriminative. Note that different choices of t can be made if matching of features at different scales is desired.

The ADF has a number of properties that make it suitable for matching in our graph-based framework.

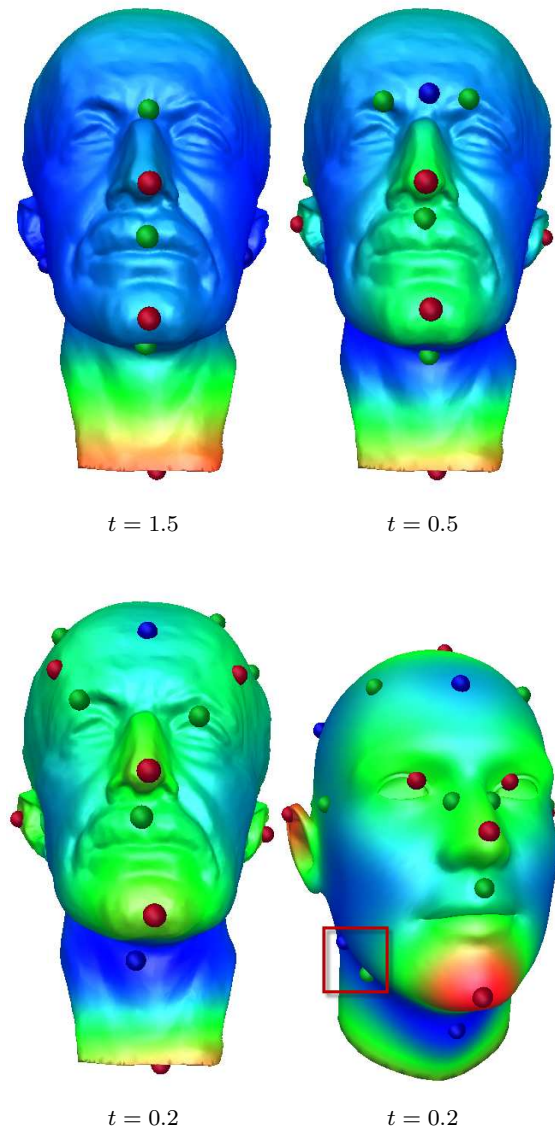


Fig. 5 ADF and sampled points with different t . Note that in the second row, points at the cheek of the right head (in rectangle) have no corresponding points on the left head.

Firstly, with the choice of a large t and small number of eigens (typically < 10), the function is generally very smooth, and hence the MSC has a compact structure with few spurious points and connectivity. Secondly, the ADF is invariant under isometric deformations, which means that if we are matching two isometric-deformed objects, their ADF should be identical, and so do the feature graphs. Thirdly, the critical points locating in features ensures our correspondences found are meaningful.

The last two properties are especially useful in sampling the candidate points. It was observed previously

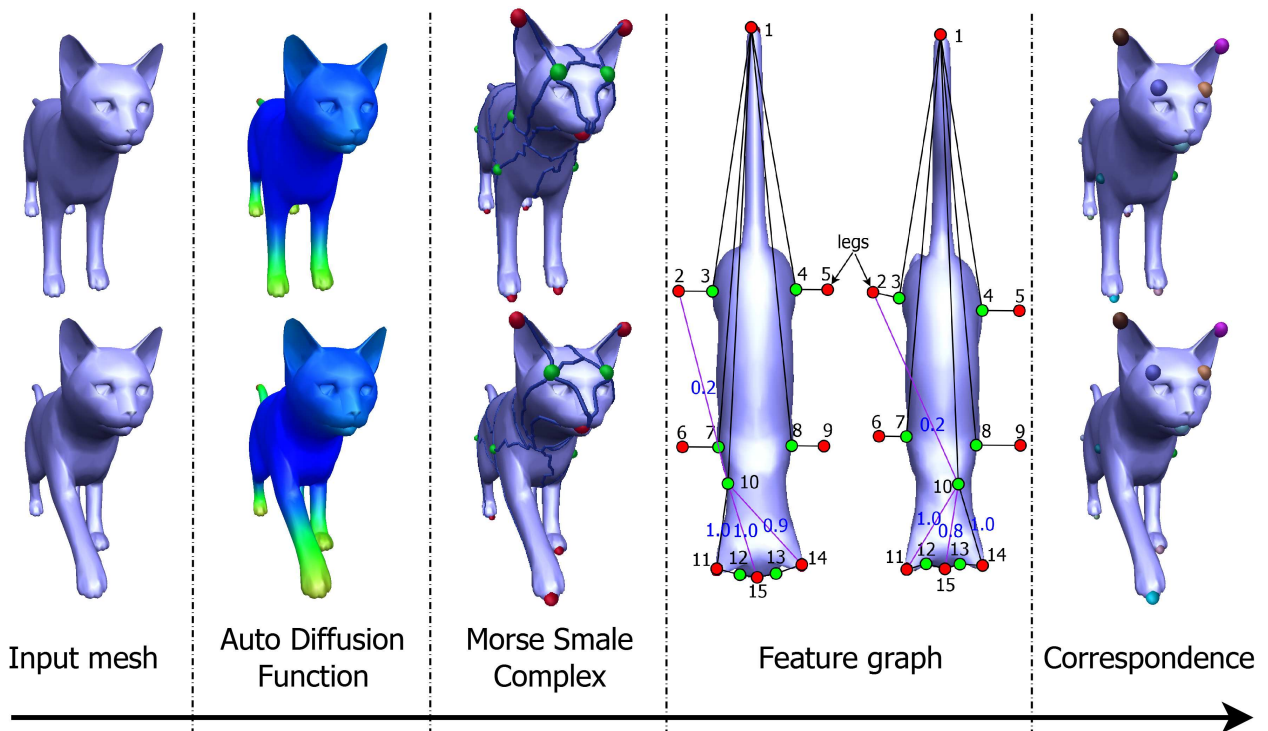


Fig. 4 The correspondence algorithm: given two input meshes, we compute ADF on the surfaces (shown in heat color) and their MSC (maxima, minima and saddles are marked in red, blue and green). All other figures use this color setting. We construct two feature graphs with nodes at the critical points and arcs from the integral lines of MSC, and perform matching of the two graphs to derive correspondences of the feature (critical) points. The purple lines are the extra added paths.

[5]) that, even under larger shape variances, the local maxima of ADF are still quite stable choices for candidates, and hence have been considered in several recent algorithms [27, 21, 28]. In our experiments, we have found that the local minima and saddles of ADF are also stable except for those at the center of the objects where the scalar values are close to the global minimum of ADF. Nevertheless, the connectivity of these critical points in the MSC remains stable in many near-isometric deformations. In addition, while maxima only appear at the extremities on a surface, the inclusion of minima and saddles gives a more uniform sample coverage, which is useful for downstream applications such as cross-parametrization. As a consequence, we consider *all* critical points of ADF as our feature points, and there is a similar approach [25], which discuss the properties of using all critical points of eigenfunctions of Laplace-Beltrami operator, which is essentially similar to ADF. We will show in the result section that the sampled points are stable under extreme affine transformations tested on the SHREC benchmark models. As an example, Figure 4 shows the ADF over two surfaces as well as their critical points. Note that the number of sampled points is not necessarily the same, especially for surfaces of different genus.

4.2 Feature Graph Construction

We construct a MSC for each surface from the sampled critical points. The critical points are identified by finding the lower and upper stars from one-ring neighborhood ([10]). Critical points are connected by *integral lines* which are maximal paths whose tangent vectors agree with the gradient. We would also apply persistence-based pruning with a persistence threshold of 0.002 if the surface is noisy. We refer readers to detailed description of MSC in [10] and to a robust computation of MSC on a triangular mesh in [4]. Note that the stable/unstable manifold ([10]), or called ascending/descending manifold in [4], does not mean that the MSC is stable or unstable. The construction of MSC is stable because of the smoothness of ADF and well-defined persistence. However there is a problem that there may have ambiguity in constructing integral lines of MSC and it would devastate our results, which will be detailed and tackled next. Given a pruned MSC, we derive a feature graph as follows:

Nodes: A node v in the graph is a critical point in the pruned MSC, and is associated with two attributes. The first attribute is the type of the critical point (maximum, minimum or saddle), denoted as $t(v)$, which will be used to enforce matching of critical points of the

same type. The second attribute is a descriptor of the local shape around the point v_i , which will be used to match features of similar appearances. While there are many such point-wise descriptors available (e.g. [26, 29]), we found that the scalar values of ADF, $f(v)$, are good enough in our examples to yield correct correspondences, because of the discriminative power of the connectivity. Note that this descriptor is a simplified HKS, and could easily be replaced by HKS or other point signatures to make a more practical system.

Arcs: The feature graph initially includes all integral lines in the MSC as its connectivity or arcs. However due to rule that a simple saddle could connect to at most two maxima, if there are multiple potential maxima for a saddle, a small perturbation could result in different connectivity in constructing the integral lines. For example, in Figure 6, saddle 2 could equally connect to maxima 3 or 5, besides maximum 1, so there are two possibilities in constructing the MSC (Figure 6 top). We resolve this ambiguity by connecting all ascending/descending paths, i.e., for every saddle, we try to find a ascending path to every maximum whose scalar value (value of ADF for this point) is larger than the scalar value of this saddle, and try to find a descending path to every minimum whose scalar value is smaller than the scalar value of this saddle, as in Figure 6 bottom. Note that this would not significantly increase the number of arcs thus still retain the discriminative power of the connectivity but will successfully resolve the ambiguity problem. For example, in Figure 6, saddle 4 would never connect to maximum 1 because the scalar value of saddle 4 is larger. Besides even if saddle 4 has smaller scalar value, it may not connect to maximum 1 due to the topology of ADF. Especially in a ADF without such arc ambiguity, no new arc will be added. This is what we get in large part of our experiments and can be found in the result in Figure 11.

In practice, the desired paths could be "cut off" due to small variances in the function. However since we are matching two surfaces with similar ADF, one integral line in one surface would imply that the corresponding integral line are highly possible to appear in the other surface. Thus if the ADF is slightly perturbed, these ascending/descending paths are highly possible to appear and a natural way of perturbation is smoothing the ADF. Specifically, we construct a series of functions $\{f_0, f_1, \dots, f_{N-1}\}$ for a given N (10 in our implementation) where f_0 is the ADF and f_{i+1} is a Gaussian filtered result of f_i . We connect a saddle v_1 with a maximum (or minimum) v_2 if there exists an ascending (or descending) path between v_1, v_2 in any of f_i , and as-

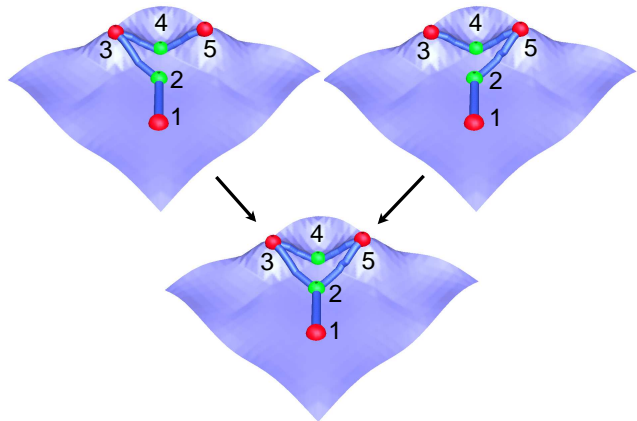


Fig. 6 Example of adding paths (the function is shown as the height of the surface). The top two figures are two possibilities of integral lines in the MSC construction. The bottom figure is the graph of connecting all ascending paths.

sociate it with an arc weight that will be used as its importance for matching:

$$w(v_1, v_2) = \frac{N - k}{N} \quad (2)$$

where k is the smallest index such that an integral line exists between v_1, v_2 in f_k . Intuitively, a smaller k means the path has higher possibility to show up after perturbation of the function. Note that the smoothing could destroy some of the critical points and make the function "drift" ([31]). We use the critical points from original ADF and use a small smoothing step (0.1 in our experiments). Also note that w is bounded between $[0, 1]$, and each existing arc from MSC has $w = 1$.

Another way of constructing the function series is using different t 's. However it suffers the same drifting and different point distribution problem as smoothing the ADF (Figure 5), or even worse. It is easier for us to control these problems by using small smoothing step, so we do not use different t for constructing ADF scales.

As a real example, Figure 4 shows part of the feature graphs on the two surfaces. Graph arcs added by finding new ascending/descending paths are colored purple. Note that the saddle number 10 is connected to different maxima on the two surfaces in the original MSC (number 11 on one surface and number 14 on the other). Using our scheme, this saddle is connected to all three maxima on the heads in both feature graphs by arcs with large arc weights, indicating that they are all likely to appear on a slightly perturbed ADF. On the other hand, a path connecting saddle 10 and maximum 2 has low weight, indicating that it is less likely to appear.

5 Graph Matching

The feature graph captures the local shape and connectivity among the feature points on a surface. The correspondence problem now reduces to one that seeks the best matching between two feature graphs, measured by compatibility of both their connectivity and attributes. In the following, we first formulate the matching problem, which is then solved efficiently using a well-known spectral technique. Due to presence of symmetry, the matching result may contain errors, which we mitigate in a final post-processing step.

5.1 Problem Statement

Consider two feature graphs $G^1 = \{V^1, E^1\}, G^2 = \{V^2, E^2\}$.

We seek a collection of *assignments* A where each assignment consists of a pair of two nodes $\{v_i^1 \in V^1, v_j^2 \in V^2\}$. To enforce one-to-one mapping, we ask nodes in all assignments to be disjoint. The goal is to find such A that maximizes a matching score defined as

$$\alpha \sum_{a \in A} H(a) + (1 - \alpha) \sum_{a, b \in A} H(a, b) \quad (3)$$

where α is a balancing weight, $H(a)$ for assignment $a = \{v_i^1, v_j^2\}$ scores the matching of the two nodes v_i^1, v_j^2 , and $H(a, b)$ for assignments $a = \{v_{i_1}^1, v_{j_1}^2\}, b = \{v_{i_2}^1, v_{j_2}^2\}$ scores the matching of the arcs (if existing) between $v_{i_1}^1, v_{i_2}^1$ in E^1 and between $v_{j_1}^2, v_{j_2}^2$ in E^2 .

More specifically, the node matching score is defined based on the types and scalar values of the nodes as:

$$H(a) = \begin{cases} 0 & \text{if } t_1(v_i^1) \neq t_2(v_j^2), \\ 1 - \left(\frac{\|f_1(v_i^1) - f_2(v_j^2)\|}{R} \right)^{0.5} & \text{otherwise.} \end{cases}$$

Here f_1, f_2 denote the two ADFs, t_1, t_2 denote the node types in the two graphs, and the normalization term $R = \max_{u \in V^1, w \in V^2} \|f_1(u) - f_2(w)\|$ ensures the range of $H(a)$ to be within $[0, 1]$. Intuitively, the higher the H , the more similar the local shapes around the two nodes are. The square root in $H(a)$ will make the differences between close signatures larger. We further prune correspondence a when $H(a) < \epsilon$, that this assignment together with its row and column will be eliminated from matrix M (see Section 5.2). We use typically $\epsilon = 0.6$ for matching isometric deformed objects for corresponding points' signatures will be very close and this is enough to prune a large number of incorrect correspondences without the risk of pruning the correct ones. For matching non-isometric deformed objects that corresponding point signatures could be of large difference, we use $\epsilon = 0.2$.

The arc matching score is defined by

$$H(a, b) = \begin{cases} 0 & \text{if } \{v_{i_1}^1, v_{i_2}^1\} \notin E^1 \\ & \text{or } \{v_{j_1}^2, v_{j_2}^2\} \notin E^2, \\ w_1(v_{i_1}^1, v_{i_2}^1) * w_2(v_{j_1}^2, v_{j_2}^2) & \text{otherwise,} \end{cases}$$

where w_1, w_2 are the arc weight functions in the two graphs. Intuitively, the higher the $H(a, b)$, the more likely it is for both node pairs $v_{i_1}^1, v_{i_2}^1$ and $v_{j_1}^2, v_{j_2}^2$ to be connected by arcs in the feature graphs on the respective surfaces. One could multiply $H(a, b)$ by the difference of arc length to respect more geometry of the surfaces. However we have found no significant improvement by doing so. Note that some points may be unmatched, e.g. due to the $\#V^1$ and $\#V^2$ are different, and they will not appear in our results of correspondences.

5.2 Spectral Matching

We solve the graph matching problem above using the spectral technique proposed by Leordeanu and Hebert [18], which avoids the combinatorial explosion inherent to the correspondence problem by exploring the spectral properties of a weighted adjacency matrix.

Specifically, consider a square matrix M where the number of rows (columns) is $|V^1| \times |V^2|$. That is, each row (column) of M represents a possible assignment between a node in V^1 and a node in V^2 . At a position in M where the row represents some assignment a and the column represents some assignment b , set the value at that position as $H(a)$ if a, b represents the same assignment, and $H(a, b)$ otherwise, where H are node and arc matching scores defined above. Note that after pruning assignments for which $H(a) < \epsilon$, the size of M becomes much smaller which allows more efficient matching. According to Perron-Frobenius theorem, the principal eigenvector x^* of M is the solution to the graph matching problem ([18]). Each value in x^* indicates how likely is an assignment to be part of a cluster of assignments with maximal inter-cluster matching score. To find assignments A , a fast, greedy selection approach is used. It starts from the highest value in x^* , adding the corresponding assignment to A , and then sets all entries in x^* whose corresponding assignments share a common node with the newly added assignment to 0. The process is repeated until no more non-zero value exists in x^* .

5.3 Order Correction

Another way of considering the better discriminative power in Figure 2 is that the incomplete graph is less

symmetric than the complete graph. Switching two points with the same connectivity (thus are symmetric) would have the same maximum in (3). Thus reducing the symmetry in the feature graph would reduce the number of wrong solutions which would have the same maximum as the correct solution and has better chance of getting the correct solution. However symmetry in object would inevitably give us symmetric feature graphs and then potentials to get symmetric parts switched results.

We propose a simple, local, iterative procedure considering orders of arcs from nodes. It is different from previous method (e.g. in [33]) that it requires less geometric information so that it is suitable for tasks of matching surfaces differing by significant deformation as we do. Specifically, for each assignment of corresponding nodes $a = \{v_i^1, v_j^2\}$ computed above, we consider the neighboring nodes of v_i^1 in G^1 whose corresponding nodes in G^2 are also neighbors of v_j^2 , and denote these two sets of corresponding neighbors as $\{N_i^1, N_j^2\}$. We mark the assignment a as *incorrect* if $|N_i^1| > 2$ and if the cyclic ordering of the arcs connecting N_i^1 with v_i^1 is different from those arcs connecting corresponding nodes in N_j^2 with v_j^2 . We maintain a queue of incorrect assignments, ordered by the decreasing value in their entry in the eigenvector of matrix M in spectral matching. At each iteration, we take the first assignment from the queue, re-assign the nodes in N_j^2 to those in N_i^1 so that the arc ordering is identical, and update the queue by adding or removing assignments whose correctness is affected by the re-assignment. To ensure termination, we require that each node in G^1 is visited at most once.

Note that the swapping procedure is limited to correct local ordering mismatches. Its success is not always guaranteed, and it cannot resolve more global errors due to large symmetric shape parts. However, we have found that the procedure works very well on the feature graphs in our tests, particularly for correcting mismatching of symmetric features (see Figure 7 for an example), which typically appear as incorrect arc ordering around isolated nodes after spectral matching.

6 Results

We now investigate the performance of the proposed approach on various examples and compare it with state-of-art methods. First, we examine in Figure 8 the variance of results under difference choices of the balance parameter α in Equation 3. Note that correct results are obtained for a large range of α , and the quality deteriorates only at extremely large values of α . This indicates that the arc connectivity of the graph (the $H(a, b)$ term in Equation 3) is a strong constraint that




Fig. 7 Correspondences on two starfishes before (first row) and after (second row) order correction.

is the main contributor to obtain the correct correspondences. We use scalar value as signature in this work only to demonstrate the discriminative power of the proposed connectivity. One could replace the point signature with more discriminative signatures and make a more practical system. For example we use HKS in Figure 9 and the direct HKS matching results are greatly improved compared to the direct matching in Figure 8 ($\alpha = 1$), with two points mismatched due to symmetry. By augmenting HKS with the proposed method, we can achieve the correct result. For HKS, we use 50 eigens and time interval $[0.1, 4]$ which was uniformly sampled 10 intervals in the logarithmically scale. Note that for the two direct matchings ($\alpha = 1$ in Figure 8 and HKS in Figure 9), we use standard Hungarian algorithm ([22]).




Fig. 8 Matching with different α values. Wrong correspondences are marked with red rectangles.




Fig. 9 Matching with HKS. In the left the points are matched directly using HKS and in the right HKS is augmented with the proposed method. Wrong correspondences are marked with red rectangle.

We next focus on comparing our connectivity-based constraints with commonly used constraints: geodesic and diffusion distances. To do so, we modify our arc matching score $H(a, b)$ in Equation 3 to be one that is based on all-pair distances:

$$H(a, b) = \exp\left(-\frac{1}{\sigma^2} \left(\frac{d_1(v_{i1}^1, v_{i2}^1) - d_2(v_{j1}^2, v_{j2}^2)}{d_1(v_{i1}^1, v_{i2}^1) + d_2(v_{j1}^2, v_{j2}^2)}\right)^2\right) \quad (4)$$

where d_1, d_2 are pairwise distances on the two surfaces in either geodesic or diffusion metric. The formula follows that of the diffusion distance in [27], where $\sigma = 0.5$ is suggested. The rest of the algorithm is left unchanged (including order correction). We compare the resulting correspondences using the three different $H(a, b)$, based on respectively our feature graph connectivity, geodesic distances and diffusion distances, on four different examples with increasing shape variance in Figure 10. Note that when the two surfaces exhibit strong isometry (e.g., the two hands), all three types of constraints work very well. As the shape variances become larger, the correspondences constrained by either geodesic or diffusion distances may contain large errors, while those constrained by our feature graph connectivity remains satisfactory.

To gain a better understanding of the stability of our feature graphs, we show in Figure 11 (top two rows) the result of matching one hand (top left) to a sequence of deformed hands from that first hand, with and without using order correction, and plot the corresponding matrices M (third row). Note that in this result there is no arc ambiguity and no new arc is added, so the arc weight is either 1 or 0. The off-diagonal part stays the same, indicating the stability of the feature graph in this deforming sequence. As a result, as the two shapes differ more and more, the correspondences deteriorate slowly with the difference of point signature. When compared with the matrices using geodesic distances as constraints instead (bottom row), we observe that M , especially the off-diagonal part, changes significantly in this sequence (see the black square and the variance matrices), indicating more drastic changes in the pairwise geodesic distances. As a result, the quality

of correspondences using geodesic constraints (fourth and fifth row) exhibit faster deterioration.

In Figure 12, we further compare our method with the Möbius voting method of Lipman and Funkhouser [19] on human models from SHREC 2007 Watertight Benchmark ([12]). These human models exhibit similar shape structures, but have relatively large shape variances between each other since they are depicting different characters. Note that the Möbius voting method produces denser correspondences while ours are sparse. However correspondences can be expanded to arbitrarily dense with an extra cross-parametrization step using these sparse but prominent feature points. We will give examples later (Figure 15). In our opinion, our correspondences are better than the results of Möbius voting in three ways: firstly, our correspondences capture prominent shape features, while in Möbius voting some prominent features are not corresponded (e.g. some feet have no correspondence). Secondly, our correspondences are consistent that all human models have (nearly) the same feature correspondences, while in Möbius voting correspondences are different for different human models. Last but not least, our correspondences are more accurate in term of point location. For example, in the first result of Möbius voting, points on the left foot of the man are corresponded to points on the left crus of the woman. Indeed there are some inconsistent and inaccurate correspondences in our results around the waist line of the human models, which will be explained next.

Obviously, our method does not always produce correct results. One common error in our correspondences are around feature points at the center of the objects or points with scalar value close to the minimum of ADF, such as those on the waistline in the human models in Figure 12. The critical points of ADF in these places tend to be much less stable than others, resulting in often large drifts of feature point locations, as visible in Figure 12. A strategy of getting high quality correspondences is removing those of both points' scalar values below a threshold. One future direction would be considering better selection of critical points as candidates for feature matching. Nevertheless, other critical points are stable, even under significant deformation. This is further exhaustively tested using SHREC'11 Correspondence Benchmark ([3], Figure 13). Note that the points on the shoulders and between the legs (all saddles) correctly appear in all human models in Figure 12 13, even under extreme affine transformations. Also, the order correction procedure cannot always correct wrong correspondences due to symmetric features, such as that between the left and right legs (Figure 13).

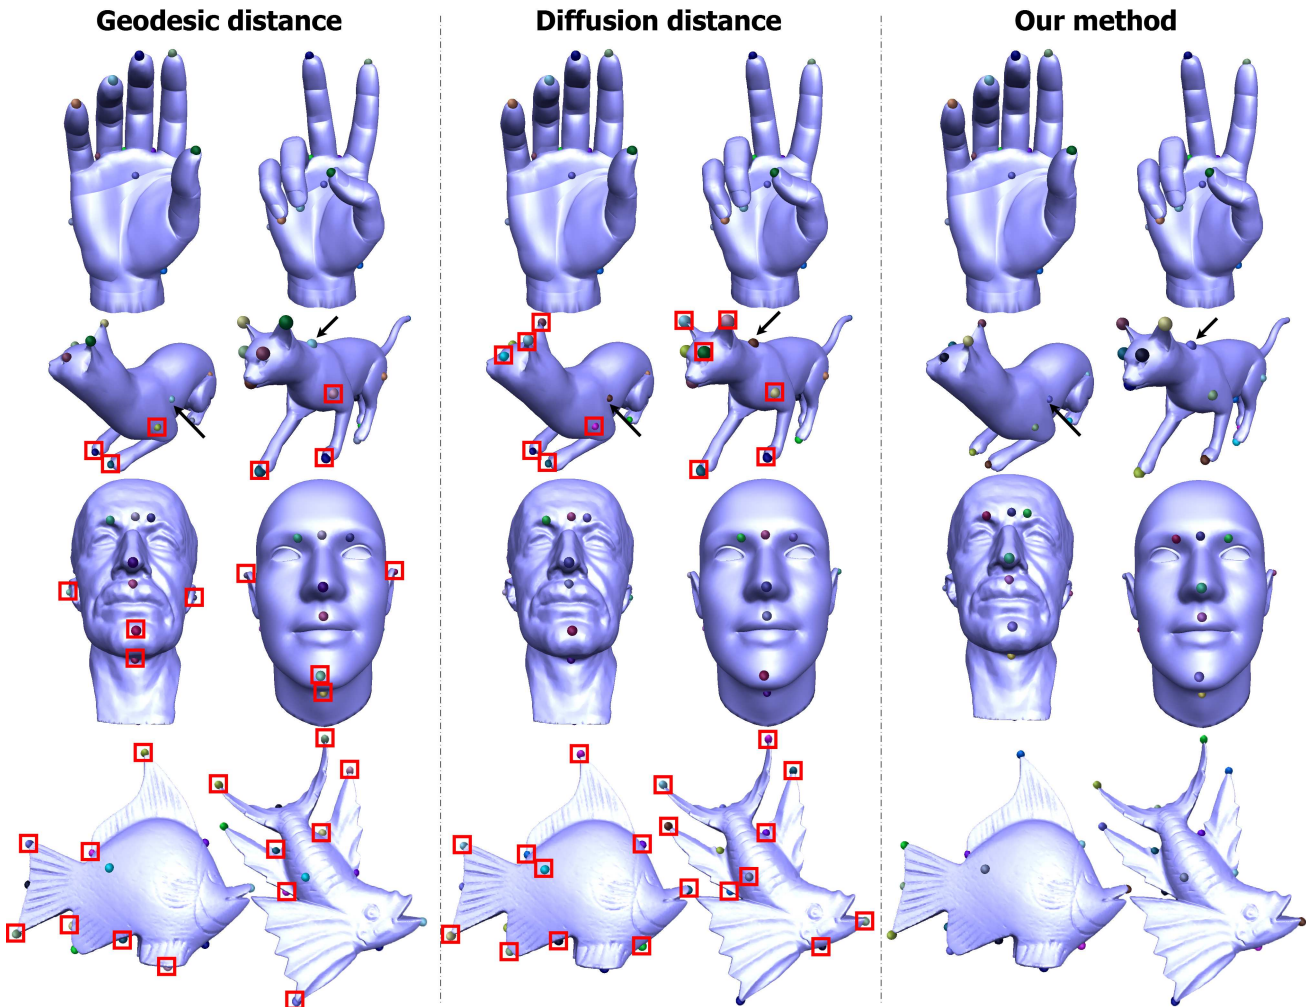


Fig. 10 Comparisons of matching results where the feature constraints are all-pair geodesic distances, all-pair diffusion distances, and the connectivity of our feature graph. The examples are ordered by decreasing isometry and shape similarity. Wrong correspondences are marked with red rectangles.

Our method works for general surfaces and is not limited to particular types of shapes or topology. For example, in Figure 14, we match both natural and CAD objects with non-zero genus. Matching between objects with different genus can also be robustly performed, as shown in Figure 1 (see accompany video for more relevant results).

Finally, to demonstrate the utility of our point correspondences, we show examples of cross-parameterization and texture transfer in Figure 15, guided by the correspondences as well as the structure of MSC. After graph matching, we prune unmatched points using MSC pruning to one of the original MSC which constrains the matched points. By using this pruning, the quadrilateral structure of the patches is retained. We then copy this MSC structure to the other MSC with the help of corresponding points ([23]). Since the two MSC graphs are topologically similar, we can only remove small amounts of non-corresponding integral lines

and add corresponding new ones. In this way, the two MSCs divides the surfaces into corresponding quadrilateral patches (shown in the first two rows of Figure 15). We then apply the parametrization method in [9] with local stiffening ([2]) to establish mapping between the patches, which further allows transferring of texture coordinates (shown in the last two rows of Figure 15). Although we are to find coarse correspondences on surface, this transferring of texture coordinates can be thought of as continuous correspondences, in which arbitrary dense point correspondences can be sampled.

Timing: Our implementation is carried out on a PC with Intel Core 2 2.80GHz and 3G RAM. The proposed method is efficient and the whole matching takes seconds to finish. Computing ADF involves a eigendecomposition of the Laplacian matrix, which typically takes 3s for a 50k triangular mesh (we only need a small amount of eigenvalues and eigenvectors). It takes less than 1s to extract MSC and another 2.5s to build




Fig. 14 Matching non-zero genus surfaces.

the feature graph. The graph matching and reordering is irrelevant to mesh resolution, and in all our experiments they take less than 0.5s. As comparisons, other methods (e.g. [33,19,32]) will typically take several minutes or even more.

7 Conclusion

In this paper, we present a simple and effective method for finding non-rigid correspondences between feature points. The core of the method is constraining the problem using a novel topological graph that is not only invariant under isometry but also observed to be stable for shapes of large variances and topological differences. The other advantage of this graph-based algorithm is its efficiency, thanks to the sparsity of the graph.

References

1. Anguelov, D., Srinivasan, P., cheung Pang, H., Koller, D.: The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In: NIPS (2004)
2. Bommes, D., Zimmer, H., Kobbelt, L.: Mixed-integer quadrangulation. In: ACM SIGGRAPH 2009 papers, SIGGRAPH '09, pp. 77:1–77:10. ACM, New York, NY, USA (2009)
3. Boyer, E., Bronstein, A.M., Bronstein, M.M., Bustos, B., Darom, T., Horaud, R., Hotz, I., Keller, Y., Keustermans, J., Kovnatsky, A., Litman, R., Reininghaus, J., Sipiran, I., Smeets, D., Suetens, P., Vandermeulen, D., Zaharescu, A., Zobel, V.: Shrec 2011: robust feature detection and description benchmark. In: Proc. Workshop on 3D Object Retrieval (3DOR'11). (2011)
4. Bremer, P.T., Edelsbrunner, H., Hamann, B., Pascucci, V.: A topological hierarchy for functions on triangulated




Fig. 15 The point and patch correspondences of two hands and two heads are shown on the first two rows. Their texture transferring results are in the following two rows.

5. surfaces. IEEE Transactions on Visualization and Computer Graphics **10**, 2004 (2004)
6. Bronstein, A., Bronstein, M., Bustos, B., Castellani, U., Crisani, M., Falcidieno, B., Guibas, L., Kokkinos, I., Murino, V., Sipiran, I., Ovsjanikovy, M., Patan, G., Spagnuolo, M., Sun, J.: Shrec 2010: robust feature detection and description benchmark. In: Eurographics 2010 Workshop on 3D Object Retrieval (3DOR'10), pp. 79–86. Eurographics Association (2010)
7. Bronstein, A.M., Bronstein, M.M., Kimmel, R., Mahmoudi, M., Sapiro, G.: A gromov-hausdorff framework

- with diffusion geometry for topologically-robust non-rigid shape matching. *Int. J. Comput. Vision* **89**, 266–286 (2010)
7. Bronstein, M.M., Kokkinos, I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In: *CVPR*, pp. 1704–1711. IEEE (2010)
 8. Coifman, R.R., Lafon, S.: Diffusion maps. *Applied and Computational Harmonic Analysis* **21**(1), 5 – 30 (2006). Diffusion Maps and Wavelets
 9. Dong, S., Bremer, P.T., Garland, M., Pascucci, V., Hart, J.C.: Spectral surface quadrangulation. *ACM Trans. Graph.* **25**, 1057–1066 (2006)
 10. Edelsbrunner, H., Harer, J., Zomorodian, A.: Hierarchical morse complexes for piecewise linear 2-manifolds. In: *Proceedings of the seventeenth annual symposium on Computational geometry, SCG '01*, pp. 70–79. ACM, New York, NY, USA (2001)
 11. Gebal, K., Baerentzen, J.A., Aanaes, H., Larsen, R.: Shape analysis using the auto diffusion function. In: *Proceedings of the Symposium on Geometry Processing, SGP '09*, pp. 1405–1413. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2009)
 12. Giorgi, D., Biasotti, S., Paraboschi, L.: Shape retrieval contest 2007: Watertight models track. <http://watertight.ge.imati.cnr.it/> (2007)
 13. Huang, Q.X., Adams, B., Wicke, M., Guibas, L.J.: Non-rigid registration under isometric deformations. In: *Proceedings of the Symposium on Geometry Processing, SGP '08*, pp. 1449–1457. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2008)
 14. Jain, V., Zhang, H.: A spectral approach to shape-based retrieval of articulated 3d models. *Comput. Aided Des.* **39**, 398–407 (2007)
 15. van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. In: *Proc. of Eurographics State-of-the-art Report*, pp. 1–22 (2010)
 16. Kim, V., Lipman, Y., Chen, X., Funkhouser, T.: Möbius transformations for global intrinsic symmetry analysis. *Computer Graphics Forum (Symposium on Geometry Processing)* **29**(5) (2010)
 17. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. *Transactions on Graphics (Proc. of SIGGRAPH 2011)* (2011)
 18. Leordeanu, M., Hebert, M.: A spectral technique for correspondence problems using pairwise constraints. In: *ICCV '05: Proceedings of the Tenth IEEE International Conference on Computer Vision*, pp. 1482–1489. IEEE Computer Society, Washington, DC, USA (2005)
 19. Lipman, Y., Funkhouser, T.: Möbius voting for surface correspondence. *ACM Trans. Graph.* **28**(3), 1–12 (2009)
 20. Milnor, J.: *Morse Theory*. Princeton Univ. Press, New Jersey (1963)
 21. Ovsjanikov, M., Mrgot, Q., Mmoli, F., Guibas, L.: One point isometric matching with the heat kernel. *Computer Graphics Forum* **29**(5), 1555–1564 (2010)
 22. Papadimitriou, C.H., Steiglitz, K.: *Combinatorial optimization: algorithms and complexity*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA (1982)
 23. Praun, E., Sweldens, W., Schröder, P.: Consistent mesh parameterizations. In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques, SIGGRAPH '01*, p. p. 179–184. ACM, New York, NY, USA (2001). DOI <http://doi.acm.org/10.1145/383259.383277>. URL <http://doi.acm.org/10.1145/383259.383277>
 24. Reuter, M.: Hierarchical shape segmentation and registration via topological features of laplace-beltrami eigenfunctions. *International Journal of Computer Vision* **89**, 287–308 (2010). 10.1007/s11263-009-0278-1
 25. Ruggeri, M.R., Patanè, G., Spagnuolo, M., Saupe, D.: Spectral-driven isometry-invariant matching of 3d shapes. *Int. J. Comput. Vision* **89**, 248–265 (2010)
 26. Rustomov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: *Proceedings of the fifth Eurographics symposium on Geometry processing*, pp. 225–233. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2007)
 27. Sharma, A., Horaud, R.P.: Shape matching based on diffusion embedding and on mutual isometric consistency. In: *Workshop on Nonrigid Shape Analysis and Deformable Image Alignment, NORDIA 2010*, June, 2010, pp. 29–36. IEEE, San Francisco, Etats-Unis (2010)
 28. Sun, J., Chen, X., Funkhouser, T.: Fuzzy geodesics and consistent sparse correspondences for deformable shapes. *Computer Graphics Forum (Symposium on Geometry Processing)* **29**(5) (2010)
 29. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: *Proceedings of the Symposium on Geometry Processing, SGP '09*, pp. 1383–1392. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2009)
 30. Tevs, A., Berner, A., Wand, M., Ihrke, I., Seidel, H.P.: Intrinsic shape matching by planned landmark sampling. In: *Eurographics*, p. to appear (2011)
 31. Weinkauff, T., Gingold, Y.I., Sorkine, O.: Topology-based smoothing of 2d scalar fields with c^1 -continuity. *Comput. Graph. Forum* **29**(3), 1221–1230 (2010)
 32. Zeng, Y., Gu, X., Samaras, D., Wang, C., Wang, Y., Paragios, N., Galen, E., de france, I.S.I.: Dense non-rigid surface registration using high-order graph matching. *CVPR* (2010)
 33. Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., van Kaick, O., Tagliasacchi, A.: Deformation-driven shape correspondence. In: *Proceedings of the Symposium on Geometry Processing, SGP '08*, pp. 1431–1439. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2008)




Fig. 11 Matching results and matrices M between one hand (top-left) and a deforming sequence of hands using our feature graph connectivity (top three rows) and geodesic distances (bottom three rows) as pairwise feature constraints. Order correction is not used in the first and fourth row and is used in the second and fifth row to give an understanding of the effect of order correction. Wrong correspondences are marked with red rectangles. A portion of visible changes in M for geodesic distance constraints is marked with black rectangle. The variances of M in which each element is $|maxvalue - minvalue|$ in the morphing sequence are shown at the first column of rows 2 and 4.




Fig. 12 Comparison of our method (second and fourth rows) with Möbius voting method (first and third rows) using human models from SHREC 2007 Watertight Benchmark.

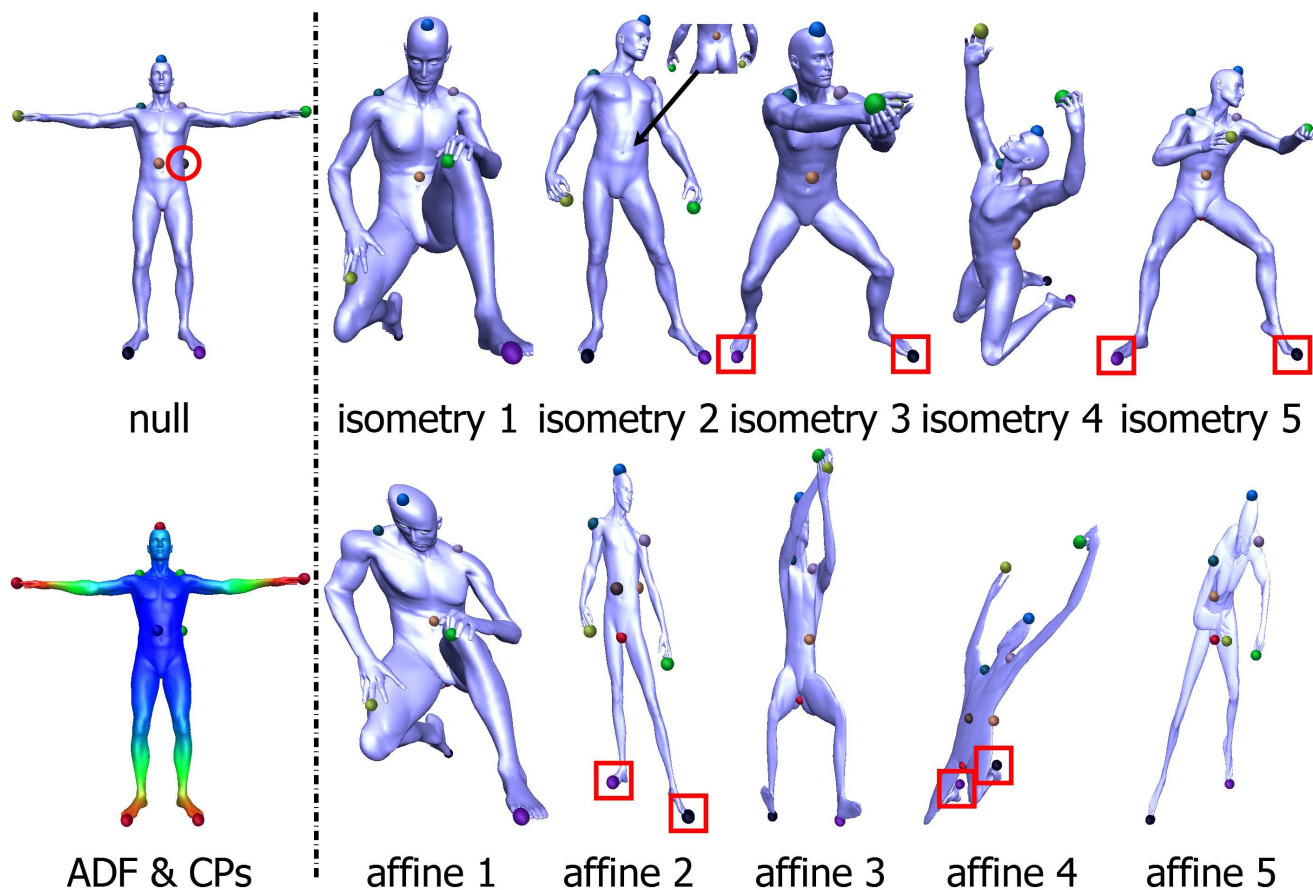


Fig. 13 Two classes of correspondences from Shrec'11 correspondence benchmark. The null shape with its' critical points (CPs) of MSC and HKS are in the left, which is matched with other ten shapes. Note that the circled points' corresponding points only appear in the affine 2-5 results. We fail to correct the order in four of the results, and these wrong correspondences are marked with rectangles.