

Bilateral Maps for Partial Matching

Oliver van Kaick, Hao Zhang and Ghassan Hamarneh

School of Computing Science, Simon Fraser University, Canada
 {ovankaic,haoz,hamarneh}@sfu.ca

Abstract

Feature-driven analysis forms the basis of many shape processing tasks, where detected feature points are characterized by local shape descriptors. Such descriptors have so far been defined to capture regions of interest centered at individual points. Using such regions to compare feature points can be problematic when performing partial shape matching, since the region of interest is typically defined as an isotropic neighborhood around a point, which does not adapt to the geometry of the shape parts. We introduce the bilateral map, a local shape descriptor whose region of interest is defined by two feature points. Compared to the classical descriptor definition using a single point, the bilateral approach exploits the use of a second point to place more constraints on the selection of the spatial context for feature analysis. This leads to a descriptor where the shape of the region of interest adapts to the context of the two points, making it more refined for shape matching. In particular, we show that our new descriptor is more effective for partial matching, since potentially extraneous regions of the models are selectively ignored owing to the adaptive nature of the bilateral map. This property also renders the bilateral map partially insensitive to topological changes. We demonstrate the effectiveness of the bilateral map for partial matching via several correspondence and retrieval experiments and evaluate the results both qualitatively and quantitatively.

1. Introduction

Many shape processing techniques are based on feature analysis, where detected feature points are characterized by local shape descriptors. Important tasks such as segmentation [Sha08], shape retrieval [TV08], correspondence [vKZHC01], and symmetry detection [XZT*09], can all be solved with the aid of local shape descriptors. A variety of such descriptors have been proposed in the literature, where a point can be represented either by a scalar property (e.g., curvature [MCH*06] or local volume [SSS*09]), or more effectively by representing the context or region of interest (ROI) around the point, as in the popular shape context descriptor [BMP02]. The latter is commonly extended to 2D manifolds by laying out a concentric grid on the surface around a point and then aggregating the geometric properties of points or faces that fall within each grid cell or bin, e.g., curvature [GGGZ05] or area [KHS10] can be summed. We call this descriptor a *geodesic map*. The main characteristic of all of these local shape descriptors is that they capture a ROI centered at a *single* feature point.

When dealing with incomplete shapes, shapes composed of a mixture of parts from multiple classes, or shapes that

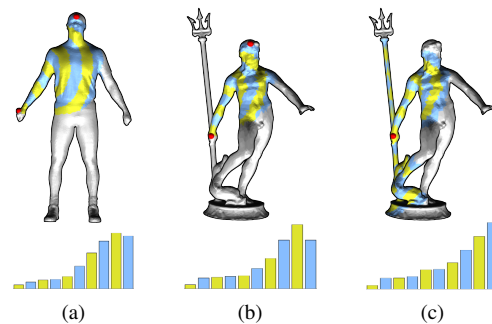


Figure 1: The bilateral map facilitates partial matching: the region captured between the head and right hand is similar in (a) and (b), and does not include extraneous parts of the Neptune model (the spear and base). In contrast, the single-point geodesic map for the right hand in (c) includes extraneous parts, even with a reduced coverage. Note that the gray portions of the shapes are not part of the regions of interest, while descriptor bins are shown with alternating colors. The content of each bin (area sum) is shown with the histograms.




Figure 2: First construction scheme for the bilateral map: (a) Starting from the shortest geodesic path between the two reference points (in red), we compute the geodesic distances from all faces in the mesh to the path (shown as a red-to-blue colormap). (b) The ROI (in yellow) is obtained by thresholding the geodesic distances.

possess significant topological variability, it becomes necessary to use descriptors that enable *partial matching*. As shown with the example in Figure 1, if we wish to match a human model to the Neptune, the extraneous parts of the Neptune (spear and base) should not be included in the context. Otherwise, the effectiveness of the descriptors is reduced, affecting the quality of retrieval and correspondence results. One solution is to first segment the shape into meaningful parts and then eliminate the extraneous parts from the ROI. However, obtaining a meaningful segmentation of a shape is a difficult problem [Sha08], as well as determining what regions are extraneous. Thus, we would rather make the descriptors independent of such a requirement. A possible straightforward solution for this problem is to assign a *scale* or *radius* parameter to the descriptor, to reduce its ROI, as shown in Figure 1 (c). However, automatically selecting the proper scale is not a trivial problem, and it is clear that this solution still has deficiencies as, no matter what radius is selected, the context of certain feature points will always include undesired portions of the models (as for points near Neptune’s right hand or near the feet).

In this paper, we propose a new type of local shape descriptor that we call the *bilateral map* (Figure 1), which is designed to circumvent these problems. Instead of defining a ROI around a single point and constraining it to a fixed radius, we compute a descriptor whose context is constrained by a *pair* of points. We propose two construction schemes to obtain such a ROI. In the first scheme, shown in Figure 2, we compute the shortest path on the surface between the pair of points, and define a ROI in the vicinity of the path according to a *region width* parameter. In the second scheme, shown in Figure 3, the ROI is obtained by thresholding the *fuzzy geodesic* [SCF10] between the two points. After constructing the ROI with one of the two schemes, we define equally spaced bins over the ROI, with bin boundaries defined by in-




Figure 3: Second construction scheme for the bilateral map: (a) We compute the fuzzy geodesic [SCF10] between the two reference points, which is a scalar field over the mesh. Notice its similarity to the distance field in Figure 2. (b) The ROI is obtained by thresholding the scalar field.

tervals of the geodesic distance emanating from one of the two points; see Figure 4. Finally, we aggregate a geometric property of all the faces or points that fall into each bin, e.g., the face areas or vertex curvatures can be summed.

It is interesting to note that in the independent parallel work of Zheng et al. [ZTZ12], point pairs have also been employed to construct a shape descriptor which is then used in applications such as intrinsic reflectional symmetry axis computation, shape correspondence, skeletonization, and segmentation. While Zheng et al. build the descriptor based on a pairwise harmonic field between two points, the construction of our bilateral map is based on the geodesic path or fuzzy geodesic between the two points.

The main observation of this paper is that it is more advantageous to define regions of interest anchored by two points instead of one point, where these regions are obtained without the need for a meaningful segmentation of the models. The bilateral map possesses several advantages over descriptors centered at a single point. First, the ROI is *adaptive*. That is, since it is constrained to lie between the two reference points, it only includes portions of the shape that capture the structural relationship between these two points. Portions of the shape that are not relevant to the reference points and that may be potentially missing in other shapes are ignored by the descriptor; contrast Figure 1 (b) to (c). Secondly, given that the ROI is adaptive, the selection of the scale of the ROI is facilitated, since the region width can be set to be proportional to the geodesic distance between the two points. In this manner, if the absolute distance between two points changes from one shape to the other, the region width is adjusted according to the distance. Finally, the bilateral approach is also partially insensitive to moderate topological changes, in comparison to other descriptors that also depend on geodesics.

Moreover, the bilateral map can have an immediate im-




Figure 4: Definition of bins for the bilateral map: the initial ROI in (a) is filtered to obtain the more restricted ROI in (b), which is then divided into equally spaced bins according to geodesic distances emanating from p , as shown in (c).

impact on existing correspondence approaches. Some of the most successful methods are based on the assumption of isometry preservation, which already requires an analysis of the compatibility between *pairs* of matches. This is the case, for example, in recent search-based methods [ZSCO*08, ACOT*10], sampling-based algorithms [TBW*11], or spectral matching methods [LH05]. Thus, our descriptors naturally fit into such approaches, since they can be added directly to the estimate of pairwise compatibility, along with other constraints such as geodesic distortion [vKZHCO11].

We demonstrate the effectiveness of the new descriptor with several experiments on shape correspondence and retrieval on complete and partial shapes, and evaluate the quality of the results in a qualitative and quantitative manner. We show that the top matches returned by the bilateral maps are generally better than those returned by descriptors centered at a single point, including the scale-invariant heat signature [BK10]. We also obtain improved results for shape retrieval when comparing to a state-of-the-art approach: the persistent heat signature of Dey et al. [DLL*10]. In addition, we show that, although the bilateral maps were designed for partial matching problems, they also display good performance when employed to match complete shapes.

2. Related work

Shape matching. A prominent problem in shape analysis is the development of means to compare the geometry of shapes, be they the full models or parts. This problem is at the heart of shape retrieval [TV08] and correspondence [vKZHCO11], which are a prerequisite of several applications, e.g., attribute transfer, shape morphing, and 3D content creation. When matching shapes with significant variability and missing data, it is important that the above applications are able to perform *partial matching*. Partial matching is difficult since, before computing the similarity of the shapes, we first need to find the common portions of the shapes [GCO06, vKZHCO11]. This requires the careful design of descriptors that are less sensitive to variations in the part composition of the models and also a mechanism to search for a partial match. The latter can be

achieved, for example, by detecting a sharp increase in the objective function when outlier regions are excluded from the match [GMGP05, ZSCO*08], or by making use of voting methods [LF09, ACOT*10].

There are also approaches that compute a partial correspondence without relying on shape descriptors, such as the method of Bronstein and Bronstein [BB08]. Their approach optimizes a type of Mumford-Shah functional, common for image segmentation, with a quasi-Newton optimization.

Local shape descriptors. A great variety of local shape descriptors have been proposed in the literature, since the most common approach to shape correspondence and retrieval is to utilize descriptors as a more suitable representation for shape comparison. Simple descriptors capture scalar properties of points, such as the integral invariants [MCH*06], or the shape diameter function [SSS*09]. Descriptors that go beyond scalar properties explicitly represent a region of interest or *context* around the points, resulting in a vector of scalar properties. The *shape context* descriptor is based on laying out a grid on this region and then counting the number of feature points that fall into the bins implied by the grid. A direct extension of this descriptor to 3D has been proposed [KPNK03], as well as a version that defines the grid relative to the surface orientation; the spin images [JH99]. Kazhdan et al. [KFR03] propose to encode such contextual descriptors in a rotationally-invariant manner by using spherical harmonics. A natural extension of shape context to manifolds is to lay out the grid on the surface of the models, and then aggregate the curvature [GGGZ05] or area [KHS10] of the surface portion that falls into the bins. Another representation encodes the region of interest with a statistical model [CCFM08].

However, these descriptors provide little flexibility to capture partial regions of the shapes, since what can be regulated is mainly the size of the context. Thus, multi-scale descriptors have been introduced, such as the multi-scale features of Li and Guskov [LG05], or integral invariants captured at different scales [MCH*06]. However, the context region captured by these descriptors is still isotropic and non-adaptive.

Gal and Cohen-Or propose a descriptor designed specifically for partial matching, where salient regions are extracted from the models and stored as a full geometry that is later matched with geometric hashing [GCO06]. Another possibility is to include part information to limit the context of the descriptors, e.g., by making use of the isophotic metric [PSH*04] or a part-aware metric [LZSCO09]. Ferreira et al. [FMA*10] extract parts from the shapes to compose a *part thesaurus* that is used for partial matching, while Itskovich and Tal [IT11] propose to combine sets of feature points and part information (segments) in the matching. Finally, a representation that has achieved considerable success in shape matching is the heat kernel signature [DLL*10, OMMG10], which can tolerate significant




Figure 5: The bilateral map is a directional descriptor: the descriptor for (p, q) is different from the one for (q, p) (the geodesic bins and histogram contents are different), preventing the confusion between p and q during the matching.

variability and missing data. This signature can also be computed in a multi-scale or scale-invariant manner [BK10].

In contrast to these works, the bilateral maps go beyond the multi-scale or fixed-size context representations and do not require part detection. The partial regions of the shapes are defined by pairs of points.

Pairwise descriptors. Pairs of points have been used before in the context of shape matching, e.g., to quantify the deviation from isometry introduced by a correspondence [vKZHC011]. However, as a shape representation, we are only aware of the work of Zheng et al. [ZTZ12] discussed in the introduction and two other works that build shape signatures based on pairs of points. In the work of Bronstein et al. [BBGO11], a histogram captures the frequency that pairs of features appear close-by on a shape, according to the pair’s diffusion distance. Being of a global nature, this signature is mainly suitable for shape retrieval. In the work of Sun et al. [SCF10], the *intersection configuration* descriptor is proposed for shape correspondence, where the amount of intersection between the fuzzy geodesics of two pairs is quantified by a single scalar value. On the other hand, the novelty of our approach is to propose a local shape descriptor that explicitly represents *the region lying in-between* the two reference points. Thus, the bilateral maps can be used for correspondence and retrieval.

3. Descriptor construction

In this section, we discuss the construction of the bilateral maps. First, given a pair of reference points (p, q) , we define a ROI that captures the context between the two points (Figures 2 and 3). The way in which we define the ROI is explained in the next heading. Next, we compute the shortest geodesic distance from each face within the ROI to the point p , defining a scalar distance field. By dividing the range of possible scalar values into b equal intervals, we are able to divide the ROI into b bins. Finally, for each bin, we aggregate a property of the faces or vertices that fall into the bin




Figure 6: Definition of the ROI for the bilateral map: (a) An initial region is defined either as the region within a distance θ from the shortest path, or by thresholding the fuzzy geodesic between p and q . (b) A filtering region (in purple at the center) is defined by intersecting the expanding fronts (in red and blue) from the two feature points. (c) The final ROI is given by the intersection of (a) and (b).

and use the result as the scalar value that represents the bin. In this paper, we simply sum the area of all the faces in the bin, although other properties can be utilized, such as the curvature or the heat kernel signature of the points. The resulting descriptor is a b -dimensional vector as shown by the histograms in Figure 1 (a) or (b). For normalization, we divide each vector entry by the sum of all entries.

Notice that our construction provides a directional descriptor. That is, the descriptor for (p, q) is different from the one computed for (q, p) . This is in fact a desired and important property to prevent the switching of any two points p and q , since it ensures that matching (p, q) to a pair (r, s) will return a different similarity than matching (q, p) to (r, s) . In the former match, p is in correspondence with r , while in the latter, q is in correspondence with r (Figure 5).

Region of interest. Our goal is to capture a ROI between the two points and avoid including extraneous portions of the models. Thus, we start the construction with an initial region obtained with one of two schemes. In the first scheme, shown in Figure 2, we create a scalar field given by the shortest distance from each face to the shortest geodesic path between the two points. All the faces with the distance below a threshold θ define our *initial* region. Before this step, we divide all the distances by the shortest distance between p and q , so that the threshold becomes relative to the path length. The parameter θ can be seen as the relative *width* of the initial region. We study the selection of this parameter in Section 4 and show that using a fixed θ is sufficiently robust.

In the second scheme, shown in Figure 3, the initial region is obtained from the *fuzzy geodesic* of two points [SCF10]. Given two points p and q , the fuzzy geodesic $G_{p,q}(x)$ is a scalar field defined on the mesh that quantifies how long is the shortest path from p to q through the diversion x , in comparison to the direct shortest geodesic from p to q :

$$G_{p,q}(x) = |\text{dist}_S(x, p) + \text{dist}_S(x, q) - \text{dist}_S(p, q)|, \quad (1)$$

where dist_S is the shortest geodesic distance between two




Figure 7: Scale selection problem on a model of two humans holding hands. Suppose that our goal is to match this model to another human with the aid of single-point descriptors. (a) For locations such as the man’s arm, we need to select smaller scales so as not to include regions of the woman’s portion. (b) If we apply such a small scale to all the points, we get less descriptive contexts. On the other hand, the bilateral maps between head and arm (c), and head and hand (d), more naturally adapt to the context of the points while being constructed with the same width parameter.

faces on the shape S . Note that we do not filter $G_{p,q}$ through a Gaussian as in [SCF10]. Finally, given a threshold τ , we create our initial region by collecting all the faces x with $G_{p,q}(x) < \tau$. However, first we normalize $G_{p,q}$ by $\text{dist}_S(p, q)$, so that the threshold τ also defines the relative width of the initial region, playing a similar role as θ . Note that using the fuzzy geodesics can create a region that is more robust when the shortest path between p and q is unstable, although there is no guarantee that the region will be connected.

After obtaining the initial region with one of the two schemes, we proceed as shown in Figure 6. We define a *filtering* region in order to exclude portions of the shape that lie beyond the reference points. Given that δ is the shortest geodesic distance between p and q , we compute the intersection of all the faces that are at most a distance of δ away from p and all the faces that are at most δ away from q . In this manner, we obtain a region that excludes the faces beyond (or “behind”) p and q , as shown in Figure 6 (b), assuming that there are no small handles or tunnels in the region. Finally, we intersect the initial and filtering regions to define the ROI for the bilateral map; see Figure 6 (c). The effect of this filtering can be seen in Figure 4, where the region highlighted in (a) is not part of the ROI after filtering.

Adaptiveness and scale selection. Our bilateral maps have several advantages over approaches centered at a single point. Firstly, the geometry of the ROI is *adaptive* and cap-




Figure 8: Partial insensitivity of the bilateral map to topological changes: when a topological shortcut exists in one of the shapes, the bilateral map capturing the region between the two feet is significantly different in (a) and (b). However, notice that many of the remaining bilateral maps are not affected by this shortcut, e.g., between the foot and left hand in (c) and (d), between the foot and points on the head, etc.

tures mainly the parts of the shape that lie between the reference points. In contrast to other descriptors, the ROI does not simply expand isotropically towards all directions; it is *anisotropically* constrained by the region between a pair of points, as shown in Figure 1 (a) and (b), and Figure 7 (c) and (d). Secondly, instead of selecting a free parameter to define the scale of the ROI, which could assume any arbitrary value, we instead select the *width* of the bilateral maps relative to the distance between the two points. This adapts the scale of the descriptor to the context of the two points on the overall shape. As shown in Figure 7, descriptors centered at a single point might require different scales depending on the location of the points on the shape, while the bilateral maps more easily adapt the size of the region while using the same width parameter.

Partial insensitivity to topological changes. Since our construction is based on the geodesic path or distance between two points, the bilateral maps are affected when the shortest path is unstable. This will happen when a *topological shortcut* exists in one of the shapes (Figure 8). Nevertheless, we observe that the paradigm of using pairs of points makes this approach partially insensitive to topological changes. Note that, when one topological shortcut exists in the shape, most of the bilateral maps will still remain intact, since the shortcut affects only the descriptors related to points in its vicinity (Figure 8). This is the case in the bilateral approach since each feature point is associated to multiple bilateral maps. Thus, if one of these maps is affected, there still remain several other maps to guide the algorithm. This property makes the *full set* of bilateral maps of a shape less sensitive (but not completely insensitive) to moderate topological changes. This is not the case in the single-point approach, since each feature point has a single associated descriptor which is then modified and unable to provide a proper match for the point.

Suitability to existing correspondence methods. Another advantage of the bilateral maps is that they can be used with existing correspondence approaches without the need for a major modification to these methods. Recent isometry-preserving methods minimize an objective similar to

$$\pi^* = \underset{\pi}{\operatorname{argmin}} \sum_{p \in S_1} \operatorname{dissim}(D_p, D_{\pi(p)}) + \alpha \sum_{p, q \in S_1} |\operatorname{dist}_{S_1}(p, q) - \operatorname{dist}_{S_2}(\pi(p), \pi(q))|, \quad (2)$$

where π^* is the optimal correspondence that we seek, p, q are points on shape S_1 , $\pi(p), \pi(q)$ are their corresponding points on shape S_2 , dissim is the dissimilarity between the descriptors D_p and $D_{\pi(p)}$, and $\operatorname{dist}_{S_i}$ is the geodesic distance on shape S_i . This objective function captures the notion that two pairs of points should be matched if the distance between the points is similar on each shape, and the descriptors of the matching points are also similar [LH05, ZSCO*08]. α regulates the balance between the geodesic inconsistency and the descriptor dissimilarity in the objective function.

Thus, to make use of our descriptors, we can simply replace the first term of the objective by

$$\sum_{p, q \in S_1} \operatorname{dissim}(D_{(p, q)}, D_{(\pi(p), \pi(q))}) \quad (3)$$

where $D_{(p, q)}$ are now the bilateral maps defined over pairs of points. Thus, the general complexity of *evaluating* (2) is not increased with this modification; only the time required to compute the descriptors is increased. Moreover, notice that the complexity of *optimizing* (2) depends on the heuristic method being used, since optimizing this functional in its general form is an NP-hard problem [vKZHCO11].

Similarly, for other applications such as shape retrieval, we will also have an increase in the time required to compute the descriptors. However, as we show in Section 4, typically a simple scheme is used to match the signatures of two shapes. Thus, in practice, we may be able to find a balance between the number of features extracted from the shapes and the quadratic number of descriptors generated.

Complexity of bilateral maps computation. When computing geodesic distances between n feature points by approximating the mesh with a discrete graph, we have a time complexity of $O(n(|E| + |V| \log |V|))$, where $|V|$ and $|E|$ are respectively the total number of vertices and edges in the mesh. Thus, the computation of the bilateral maps is dominated by this complexity. In practice, computing the 2,450 bilateral maps between 50 feature points on a mesh of 20K faces takes 10 minutes with unoptimized MATLAB code. In contrast, computing the corresponding 50 single-point geodesic maps takes 1 minute with the same base code.

4. Experiments and results

In this section, we demonstrate the effectiveness of the bilateral maps in the contexts of shape correspondence

and retrieval. For our initial experiments, we use the first scheme proposed to construct the region of interest between two points, where the region is grown around the shortest geodesic path between the two points. We use a single set of descriptor parameters: number of bins $b = 20$, and $\theta = 0.30$. We also study the impact that different values of these parameters have on the correspondence results, as well as the second construction scheme of the initial region.

4.1. Shape correspondence

Although a variety of methods that make use of local shape descriptors for correspondence have been proposed in the literature [vKZHCO11], we chose to perform a *raw* evaluation of the bilateral maps, i.e., pairwise matches are obtained from the descriptors without any intermediate optimization or filtering steps. In this manner, the effectiveness of the bilateral maps can be isolated and evaluated directly.

However, given that the bilateral maps are defined between pairs of points, we still need a mechanism to extract a point-to-point correspondence from the descriptors defined between pairs of points. Following the spirit of a direct evaluation, we use a simple voting algorithm for this step. Given two meshes \mathcal{M}_1 and \mathcal{M}_2 , for each bilateral map on \mathcal{M}_1 , we find the most similar map on \mathcal{M}_2 , and place a vote on the two pairwise matches implied by the two descriptors. We utilize the ℓ_1 -norm to measure the descriptor dissimilarity. At the end of the voting, we select for each feature point in \mathcal{M}_1 , the matching point in \mathcal{M}_2 with the largest number of votes. We also constrain the matching to be one-to-one by retaining for each point in \mathcal{M}_2 , the match to the corresponding point in \mathcal{M}_1 with the most votes. In our experiments, we sample 50 feature points uniformly across each surface and compute the bilateral maps for all the pairs of points. We display only the top 50% matches, since correspondence algorithms will typically explore only the top matches.

Results. In Figure 9, we show visual examples of correspondences computed with our bilateral maps. Notice how the descriptors are able to provide meaningful matches for shapes with missing parts or even differing topology. In (a), the method is applied to a pair of complete shapes, to show that the bilateral maps are also suitable for computing full correspondences between shapes in different poses. The examples in (b)-(h) present cases requiring partial matching. We see that in the simple case where a shape is matched to a version of itself cut in half, shown in (b), the descriptors are able to provide the correct matches between all the selected feature points. In (c)-(e), we see that the bilateral maps are effective for matching incomplete shapes, with missing parts. Finally, in (f)-(h), we show the effectiveness of the bilateral maps for matching hybrid shapes or shapes with topological differences. Notice that, since these results were obtained without any form of optimization, incorrect matches are also returned. However, note that each prominent shape part is not missed, receiving at least one correct match.

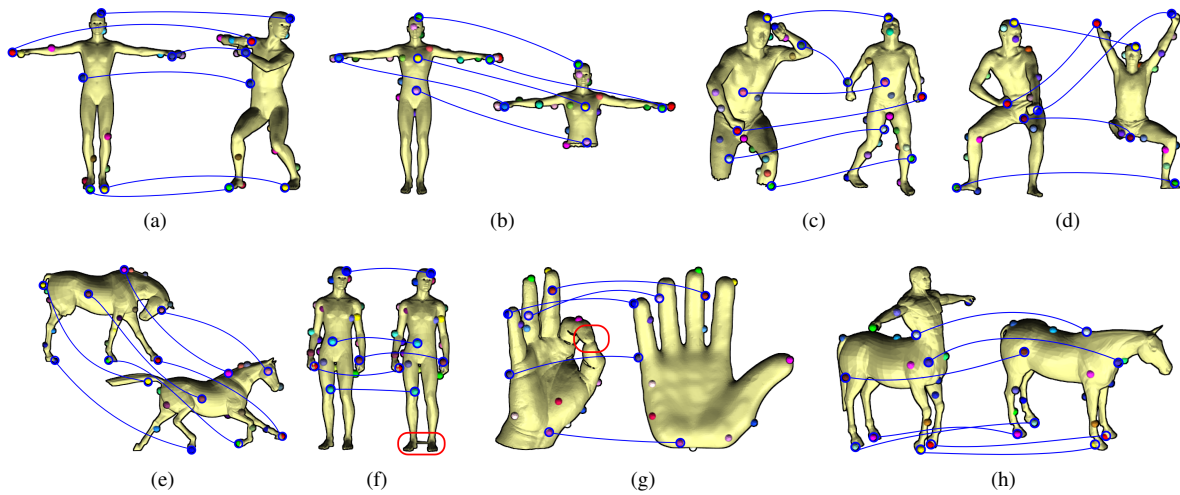


Figure 9: Correspondence results obtained with our bilateral maps on a set of example pairs. Corresponding points are shown with matching colors, and we connect some interesting matches with the blue curves. Notice that our descriptors enable to find meaningful matches between: (a) two complete shapes in different poses, (b)-(e) shapes with missing parts (missing legs, hand and tail), and (f)-(h) shapes that possess different topology (circled in red in (f) and (g)) or which include extraneous parts (the human part of the centaur in (h)). The same descriptor width θ was used in all the examples.

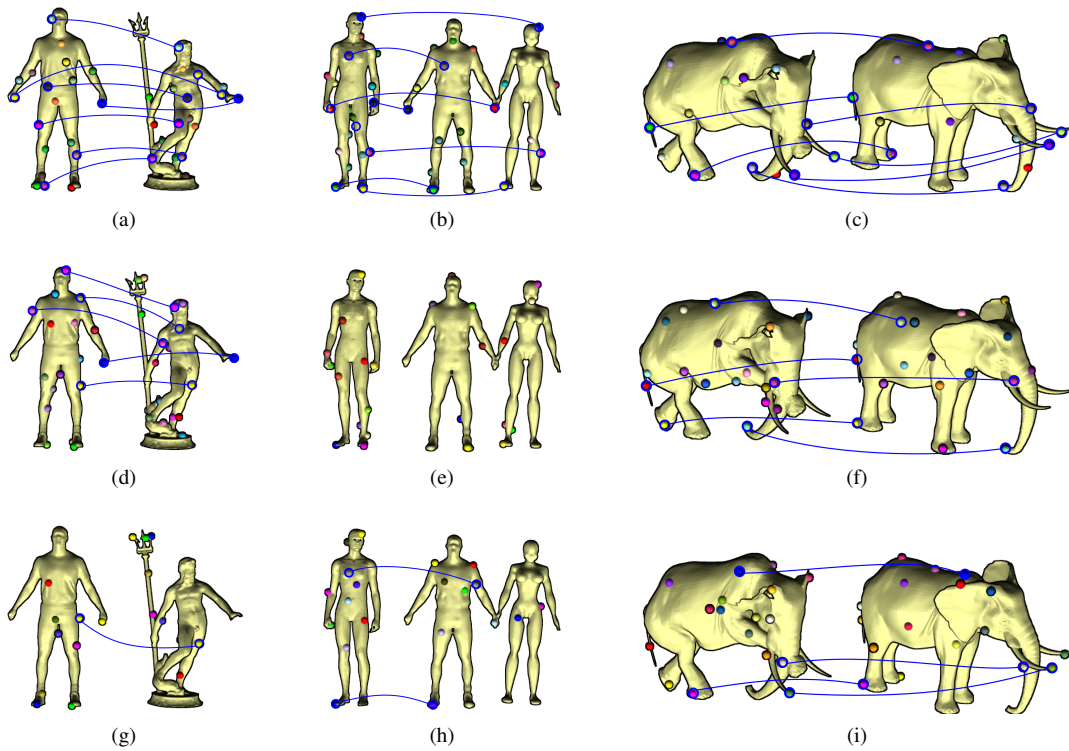


Figure 10: Correspondence results obtained with the bilateral maps in (a)-(c), compared to the results obtained with the single-point geodesic maps in (d)-(f), and the scale-invariant heat signature in (g)-(i). Note that the results obtained with the bilateral maps are more meaningful than those obtained with the single-point descriptors. Meaningful matches are highlighted in blue.

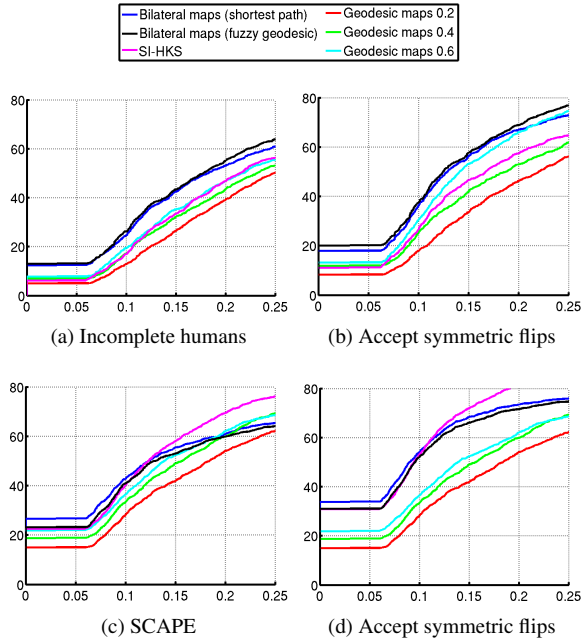


Figure 11: Quantitative evaluation of correspondences computed with the bilateral maps and single-point descriptors on incomplete (a)-(b) and complete humans (c)-(d). The x -axis denotes the correspondence error in terms of geodesic distances, while the y -axis denotes the percentage of correspondences with an error below a given value. Note the better performance of the bilateral maps on incomplete shapes, and competitive performance on the SCAPE dataset.

Comparison to single-point descriptors. Figure 10 shows additional results obtained with the bilateral maps as well as a comparison to traditional single-point descriptors: the geodesic map and the scale-invariant heat kernel signature [BK10]. The geodesic maps use the same type of intrinsic binning as the bilateral maps, and we compute these descriptors at five different scales. The extent of the descriptors is set to 20%, 40%, 60%, 80%, and 100% of the longest geodesic path on each shape. Each descriptor is composed of 20 concentric bins, where we also sum the area of the faces that fall within each bin. Next, we compute correspondences with each scale and select the best result for each example in Figure 10 (d)-(f). By using these geodesic maps that have a similar setup to the bilateral maps, we are able to compare directly the advantage of using pairs of points over single points. The heat kernel signatures are computed according to the construction in [BK10], and we use the first 10 coefficients for the signature, as suggested in that work. To establish the correspondences, we select, for each feature point, the matching point with the greatest similarity.

In Figure 10 (a), we show the matches obtained for our motivating example of the human vs. Neptune. Cor-

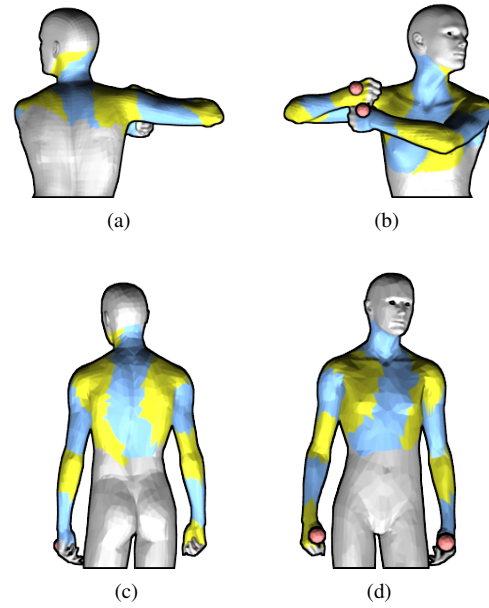


Figure 12: Context regions built with fuzzy geodesics. Here we show the bilateral maps built with the same parameters between two points on the hands, for two different shapes (front and back). Note how the resulting context regions do not correspond across the two shapes. Since more shortest paths that differ by a given ϵ tend to cross the front of the shape in (a) and (b), the region within a certain threshold is significantly different from that in (c) and (d), where similar shortest paths divide equally between front and back.

rect matches are obtained between the prominent parts of the models, including points close to topological differences (legs), while not many matches which connect extraneous regions (the spear and base) are returned. However, the right arm of the human is matched to Neptune’s left arm. Due to the intrinsic nature of the descriptor, symmetric parts are not always distinguished. In (b), we see a human matched to a single manifold composed of two humans holding hands, where all the correspondences are meaningful up to symmetry switching. Note that since we are performing a raw comparison, points on the human on the left can be matched both to the man and the woman portions of the model on the right. Finally, in (c), the left elephant has a different topology from the one on the right (its head, trunk, and tusks are connected to a front leg), but correct matches are still obtained between the trunks, tusks, and the bodies of the elephants.

Moreover, by contrasting the results of the bilateral maps in Figure 10 (a)-(c) to the results of the single-point geodesic map in (d)-(f), we see that the bilateral maps return more meaningful matches, while using the same descriptor width parameter. When selecting the proper scale, the single-point descriptors are able to match important features such as the

head and arms of the humans, and the body of the elephants. However, the existence of extraneous regions and topological changes have a stronger impact on the single-point descriptors than on the bilateral maps. Incorrect matches are returned, for example, between the legs of Neptune and the body of the human. The single-point approach completely fails in (d), since the geodesic maps of the two shapes are significantly different, even at the small scale of 20%. In comparison to the scale-invariant heat signature in (g)-(i), we see that this signature is effective for matching shapes with scale variations or small topology changes (like the elephants), but fails when the models have significant extraneous portions (like on the two humans holding hands).

Topological changes. One advantage of the bilateral maps is that, if two shapes possess moderate topological differences, we can still obtain a meaningful correspondence between the two. As discussed in Section 3, although some of the descriptors are modified when one of the shapes suffers topological changes, enough of the multiple bilateral maps associated to each feature still remain intact if only a few topological shortcuts exist. The partial insensitivity of the bilateral maps is demonstrated in the examples in Figure 9 (f) and (g), as well as in the matching of a human to Neptune in Figure 10 (a) and the elephants in (c), where more meaningful matches are returned in comparison to the single-point descriptors.

Quantitative evaluation. We also compare the bilateral maps to the single-point geodesic maps and scale-invariant heat signature in a quantitative manner on datasets of partial and complete shapes. For partial shapes, we selected 13 incomplete humans from the dataset for partial shape retrieval of Dey et al. [DLL*10]. These humans have missing parts or missing data in the form of multiple holes on the meshes. Next, we computed correspondences between each human and 8 other randomly selected humans. For evaluating the bilateral maps on complete shapes, we use the SCAPE dataset with 71 shapes [ASP*04], where we computed a correspondence between each model and another randomly-selected model in the set. To compute the correspondences, we also perform a raw matching on sets of 50 sample points.




Figure 13: Quantitative evaluation of correspondences computed with the bilateral maps with different numbers of bins and increasing descriptor width θ . Note that the best performance on this dataset is with 20 bins and θ higher than 0.3.

To evaluate the results, the SCAPE dataset has ground-

truth correspondences available. For the partial dataset, we first created a ground-truth for each shape. We selected 36 consistent landmark points on each shape, as in the benchmark of Kim et al. [KLF11]. In the case of features appearing in missing parts, we placed the ground-truth at the location nearest to the landmark. For example, if a human is missing its left leg from the knee downwards, we placed the landmark corresponding to the left toe at the location of the left knee. Since the descriptors are computed on automatic features, we then derive a ground-truth for the features from the ground-truth of the landmark points. Each feature point is represented by a vector of geodesic distances to the landmarks. By finding the best match for such vectors, we establish the ground-truth for the feature points.

Once we are given a correspondence π computed with one of the descriptors, we compute its error in relation to the ground-truth correspondence π_{gt} by adding up the geodesic distance from each selected point to the ground-truth point,

$$\text{Error}(\pi, \pi_{\text{gt}}) = \sum_{p \in \mathcal{M}_1} d_{\mathcal{M}_2}(\pi(p), \pi_{\text{gt}}(p)). \quad (4)$$

The geodesic distances $d_{\mathcal{M}_2}$ are normalized by $\sqrt{\text{Area}(\mathcal{M}_2)}$ [KLF11].

Figure 11 shows the results of this experiment. Each curve denotes the percentage of correspondences that have an error below a given geodesic distance. We also show the results when symmetric flips are accepted in the correspondence. It can be clearly seen that the bilateral maps have a better performance on the incomplete shapes, returning correspondences with lower errors. Note that, for this dataset, the geodesic maps with a larger scale of 60% gave the best results for single-point descriptors, since this larger scale better captures the human structure. For the complete shapes, the single-point descriptors are expected to perform well. However, the bilateral maps still remain as a competitive option, returning lower error correspondences on the first 40% matches. Note that correspondence algorithms will typically make use of the top most similar matches for guidance. Thus, it is essential that the top matches are of higher quality.

Figure 11 also compares the two possible constructions of the initial region of the bilateral maps: using the shortest path between the two reference points or their fuzzy geodesic. For the fuzzy geodesic construction, we chose a parameter τ that creates regions of the same size to the selected θ . We see that the two options are equivalent to each other in practice. Our observation on this result is that, if the shortest path is stable, the regions defined by both constructions are practically the same. However, in the case of instability, the more stable fuzzy geodesics do not necessarily guarantee that the context regions will correspond across the shapes. E.g., in one shape the geodesics could be more unstable than in the other, resulting in significantly different context regions. One such example is presented in Figure 12.

Figure 13 investigates the effect of using different num-

bers of bins and descriptor widths θ for the construction of the bilateral maps with the shortest paths. We evaluate the different parameters on the set of incomplete shapes. We see that, for this dataset, any θ value equal or higher than 0.30 gives good results, while increasing the number of bins only slightly improves the quality of the correspondences. We conjecture that this is the case since, although the shapes in the dataset are incomplete, they still possess large common portions, and larger contexts then capture more shape information. On the other hand, lower values for θ clearly do not capture much context around the pair of points and are less effective. We select 0.30 as the parameter in our experiments since then we guarantee that less extraneous parts will be captured when the shapes have more variations.

Limitations. The limitations of the bilateral maps can be observed in Figure 9. As common with intrinsic descriptors [KLF11], when the shapes have strong intrinsic symmetries, e.g., left- and right-halves of a human, the bilateral maps are unable to tell the symmetries apart. Also, since the descriptors are built on geodesic distances and the area of the shapes, significant distortion from isometry on the shape parts can provide poor matches, as shown in Figure 14 (a). Finally, as seen in Figure 14 (b), if several topological shortcuts exist on a shape (between arms and legs on the human on the right), the bilateral maps become significantly dissimilar and do not provide good matches for the parts involved in the shortcuts.

Another point for practical consideration is that the bilateral maps can only capture as much detail as present on the input meshes. That is, if two feature points are close-by and there are not many details in their context region (e.g., between the thumb and index finger in a low-resolution model of a complete human), then their corresponding bilateral map will not be very informative. Thus, the strength of the bilateral maps is more on what we term “medium-scale” partial matching, i.e., where the partial portions of the shapes being matched represent a significant fraction of the shapes (at least 20% or 30% of the shapes), as in the examples shown throughout this section.




Figure 14: Examples illustrating the limitations of the approach. The performance of the bilateral maps is decreased in the presence of: (a) Significant distortion from isometry. (b) Topological shortcuts affecting several parts in a shape.

4.2. Shape retrieval.

Dataset and methodology. We show results of using our bilateral maps for shape retrieval and also compare to the method of Dey et al. [DLL*10]. Their method utilizes the persistent heat kernel signature (PHS) to detect interest points on the shapes and also to define shape signatures for retrieval. Their work is the first to evaluate results specifically on incomplete shapes and also to build a dataset for this purpose. We perform the comparison by evaluating the bilateral maps on Dey et al’s dataset. The dataset is composed of 300 shapes organized into 21 classes (humans, horses, chairs, etc.). Its query set is composed of 18 complete shapes and 32 incomplete shapes, while the target set contains 197 complete and 101 incomplete shapes.

To perform the retrieval, we also utilize a procedure similar to that used by Dey et al. First, we compute the bilateral maps for 20 interest points uniformly sampled on each shape. We take the resulting 380 descriptors to create the signature of each shape. Next, during the retrieval phase, the matching score between two shapes \mathcal{M}_1 and \mathcal{M}_2 is given by the expression [DLL*10]

$$\sum_{f_1 \in F_1} \min_{f_2 \in F_2} \|f_1 - f_2\|_1 + \sum_{f_2 \in F_2} \min_{f_1 \in F_1} \|f_1 - f_2\|_1, \quad (5)$$

where F_1 and F_2 are the sets of signatures for \mathcal{M}_1 and \mathcal{M}_2 , respectively, $f_1 \in F_1$ and $f_2 \in F_2$ denote single descriptor vectors in the signature set, and $\|\dots\|_1$ is the ℓ_1 -norm. This score is simple and efficient to compute for large datasets.

For each query, the target shapes are ordered according to the score in (5) and the results are evaluated in terms of the top-3 and top-5 hit rates. Basically, given a query shape, we have a top- k hit if a shape from the same class is retrieved within the top k matches. For a class with N shapes, the top- k hit rate is the percentage of top- k hits with respect to N .

Results. Table 1 shows the retrieval hit rates for our descriptors, compared to the method of Dey et al. [DLL*10] (PHS) and the Eigenvalue descriptor method of Jain and Zhang [JZ07] (EVD). Note that, in the case of the bilateral maps, we are using a quadratic number of descriptors in comparison to a linear number for PHS and EVD. We see from the results that, although PHS and EVD have different performances on complete and incomplete shapes, our descriptors provide better results than these two methods on both types of shapes. Thus, although the bilateral maps were designed for partial matching, their performance does not drop when matching and retrieving complete shapes.

5. Conclusion

The key message of this paper is to show that it is worthwhile to consider two points instead of one point in defining regions of interest for local shape descriptors. To demonstrate this idea, we proposed the *bilateral map*, which anisotropically adapts its shape to the region comprised between two

Table 1: Retrieval results on a dataset of complete and incomplete shapes. Each table entry shows top-3 and top-5 hit rates for our bilateral maps, PHS, and EVD.

Queries	Ours	PHS	EVD
32 incompl.	91% / 94%	88% / 91%	62% / 62%
18 compl.	100% / 100%	78% / 83%	100% / 100%
50 total	94% / 96%	84% / 88%	76% / 76%

points. We showed that the bilateral map offers a promising alternative to the classic descriptor definition for tasks such as shape correspondence and retrieval. The bilateral map was designed to deal better with the partial matching problem, but we showed that this does not hinder its performance on the full matching problem.

We have shown results using the bilateral maps within a simple correspondence framework, to more easily evaluate the effectiveness of the descriptors. However, these descriptors can potentially help improve the results of more elaborate methods based on the feature analysis approach. For example, they can be readily used within recent search-based algorithms [GMGP05, ZSCO*08, ACOT*10], sampling-based algorithms [TBW*11], or even to guide transformation-based [LF09, KLF11] or registration-based methods [BB08] more rapidly towards a good solution. Moreover, we evaluated the bilateral maps by collecting the area of the faces that fall within each bin. However, the descriptor can be potentially used with other geometric properties, such as curvature [GGGZ05], the shape diameter function [SSS*09], or even with more sophisticated scalar fields, such as the heat diffusion field [DLL*10].

In the future, we plan to investigate alternative constructions of the proposed descriptors. For example, we can additionally split the bins according to the distance from the geodesic path, obtaining a 2D instead of 1D bin layout. Moreover, we can attain a higher-order construction by considering three points on a surface, which define a common region. A grid can then be defined by dividing the region along every pair of points, resulting in another 2D bin layout. Such a generalization lifts the need to specify a threshold θ for the descriptor, but increases the matching complexity.

Finally, another important topic for future investigation is that of *feature selection*, which is a difficult problem when the goal is to capture features that appear consistently across two different shapes. We evaluated the proposed bilateral maps with feature points uniformly sampled across the surfaces. However, an interesting question is whether there is a feature selection approach that is more suitable for the new descriptors. Previous works proposed to use sampling schemes based on extremity selection [ZSCO*08], saliency [CCFM08] or information criteria such as entropy [TBW*11], but the particular choice of method is highly correlated with the category of shapes be-

ing considered. Perhaps the best way of selecting feature points on surfaces is to use an approach similar to that of Tevs et al. [TBW*11], where the feature selection is part of the correspondence computation and, potentially, could even be made part of the descriptor construction.

References

- [ACOT*10] AU O. K.-C., COHEN-OR D., TAI C.-L., FU H., ZHENG Y.: Electors voting for fast automatic shape correspondence. *Computer Graphics Forum (Proc. EUROGRAPHICS)* 29, 2 (2010), 645–654.
- [ASP*04] ANGUELOV D., SRINIVASAN P., PANG H.-C., KOLLER D., THRUN S., DAVIS J.: The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Proc. of the NIPS* (Vancouver, BC, Canada, 2004).
- [BB08] BRONSTEIN A. M., BRONSTEIN M. M.: Regularized partial matching of rigid shapes. In *Proc. Euro. Conf. on Comp. Vis. (ECCV)* (Marseille, France, 2008), pp. 143–154.
- [BGO11] BRONSTEIN A. M., BRONSTEIN M. M., GUIBAS L. J., OVSJANIKOV M.: Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Trans. on Graphics* 30, 1 (2011), 1–20.
- [BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. IEEE Conf. on CVPR* (San Francisco, CA, USA, 2010), pp. 1704–1711.
- [BMP02] BELONGIE S., MALIK J., PUZICHA J.: Shape matching and object recognition using shape context. *IEEE PAMI* 24, 4 (2002), 509–522.
- [CCFM08] CASTELLANI U., CRISTANI M., FANTONI S., MURINO V.: Sparse points matching by combining 3D mesh saliency with statistical descriptors. *Computer Graphics Forum (Proc. EUROGRAPHICS)* 27, 2 (2008), 643–652.
- [DLL*10] DEY T. K., LI K., LUO C., RANJAN P., SAFA I., WANG Y.: Persistent heat signature for pose-oblivious matching of incomplete models. *Computer Graphics Forum (Proc. SGP)* 29, 5 (2010), 1545–1554.
- [FMA*10] FERREIRA A., MARINI S., ATTENE M., FONSECA M., SPAGNUOLO M., JORGE J., FALCIDIENO B.: Thesaurus-based 3D object retrieval with part-in-whole matching. *Int. J. Comput. Vision* 89, 2 (2010), 327–347.
- [GCO06] GAL R., COHEN-OR D.: Salient geometric features for partial shape matching and similarity. *ACM Trans. on Graphics* 25, 1 (2006), 130–150.
- [GGGZ05] GATZKE T., GRIMM C., GARLAND M., ZELINKA S.: Curvature maps for local shape comparison. In *Proc. Conf. on Shape Modeling and Applications* (Cambridge, MA, USA, 2005), pp. 244–253.
- [GMGP05] GELFAND N., MITRA N. J., GUIBAS L. J., POTTMANN H.: Robust global registration. In *Proc. Symp. on Geom. Processing (SGP)* (Vienna, Austria, 2005), pp. 197–206.
- [IT11] ITSKOVICH A., TAL A.: Surface partial matching and application to archaeology. *Computers & Graphics* 35, 2 (2011), 334–341.
- [JH99] JOHNSON A., HEBERT M.: Using spin-images for efficient multiple model recognition in cluttered 3D scenes. *IEEE PAMI* 21, 5 (1999), 433–449.
- [JZ07] JAIN V., ZHANG H.: A spectral approach to shape-based retrieval of articulated 3D models. *Computer-Aided Design* 39, 5 (2007), 398–407.

- [KFR03] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Proc. Symp. on Geom. Processing (SGP)* (Aachen, Germany, 2003), pp. 156–164.
- [KHS10] KALOGERAKIS E., HERTZMANN A., SINGH K.: Learning 3D mesh segmentation and labeling. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 29, 3 (2010).
- [KLF11] KIM V., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 30, 4 (2011).
- [KPNK03] KÖRTGEN M., PARK G.-J., NOVOTNI M., KLEIN R.: 3D shape matching with 3D shape contexts. In *Proc. 7th Central European Seminar on Computer Graphics* (Budmerice, Slovakia, 2003).
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 28, 3 (2009), 1–12.
- [LG05] LI X., GUSKOV I.: Multi-scale features for approximate alignment of point-based surfaces. In *Proc. Symp. on Geom. Processing (SGP)* (Vienna, Austria, 2005).
- [LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *Proc. Int. Conf. on Comp. Vis. (ICCV)* (Beijing, China, 2005), vol. 2, pp. 1482–1489.
- [LZSCO09] LIU R., ZHANG H., SHAMIR A., COHEN-OR D.: A part-aware surface metric for shape analysis. *Computer Graphics Forum (Proc. EUROGRAPHICS)* 28, 2 (2009), 397–406.
- [MCH*06] MANAY S., CREMERS D., HONG B.-W., YEZZI A. J., SOATTO S.: Integral invariants for shape matching. *IEEE PAMI* 28, 10 (2006), 1602–1618.
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *Computer Graphics Forum (Proc. SGP)* 29, 5 (2010).
- [PSH*04] POTTMANN H., STEINER T., HOFER M., HAIDER C., HANBURY A.: The isophotic metric and its application to feature sensitive morphology on surfaces. In *Proc. Euro. Conf. on Comp. Vis. (ECCV)* (Prague, Czech Republic, 2004), pp. 560–572.
- [SCF10] SUN J., CHEN X., FUNKHOUSER T.: Fuzzy geodesics and consistent sparse correspondences for deformable shapes. *Computer Graphics Forum (Proc. SGP)* 29, 5 (2010).
- [Sha08] SHAMIR A.: A survey on mesh segmentation techniques. *Computer Graphics Forum* 27, 6 (2008), 1539–1556.
- [SSS*09] SHAPIRA L., SHALOM S., SHAMIR A., COHEN-OR D., ZHANG H.: Contextual part analogies in 3D objects. *Int. J. Comput. Vision* 89, 2–3 (2009), 309–326.
- [TBW*11] TEVS A., BERNER A., WAND M., IHRKE I., SEIDEL H.-P.: Intrinsic shape matching by planned landmark sampling. *Computer Graphics Forum (Proc. EUROGRAPHICS)* 30 (2011), 543–552.
- [TV08] TANGELDER J. W. H., VELTKAMP R. C.: A survey of content based 3D shape retrieval methods. *Multimedia Tools and Applications* 39, 3 (2008), 441–471.
- [vKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707.
- [XZT*09] XU K., ZHANG H., TAGLIASACCHI A., LIU L., LI G., MENG M., XIONG Y.: Partial intrinsic reflectional symmetry of 3D shapes. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)* 28, 5 (2009).
- [ZSCO*08] ZHANG H., SHEFFER A., COHEN-OR D., ZHOU Q., VAN KAICK O., TAGLIASACCHI A.: Deformation-driven shape correspondence. *Computer Graphics Forum (Proc. SGP)* 27, 5 (2008), 1431–1439.
- [ZTZX12] ZHENG Y., TAI C.-L., ZHANG E., XU P.: Pairwise harmonics for shape analysis. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012).