

Geodesic Iso-Curve Signature

Anne Gehre

David Bommes

Leif Kobbelt

Visual Computing Institute, RWTH Aachen University




Figure 1: Illustration of the geodesic iso-curve signature: For the computation of the signature at a surface point x geodesic circles are computed around x for a set of geodesic radii. The descriptor captures the lengths of the isocontours of these circles. To the left three shapes from the FAUST shape repository with isocontours at different radii are depicted. Further, the respective surface signatures (middle) for a point located on the palm of their hands are shown. On the right the signature distance of a point located on the palm of the hand of the female shape to all other points (across shapes) are shown. The red color indicates low distance values.

Abstract

During the last decade a set of surface descriptors have been presented describing local surface features. Recent approaches [COO15] have shown that augmenting local descriptors with topological information improves the correspondence and segmentation quality. In this paper we build upon the work of Tevs et al. [TBW*11] and Sun and Abidi [SA01] by presenting a surface descriptor which captures both local surface properties and topological features of 3D objects. We present experiments on shape repositories that are provided with ground-truth correspondences (FAUST, SCAPE, TOSCA) which show that this descriptor outperforms current local surface descriptors.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Computational Geometry and Object Modeling—Line and curve generation

Keywords: shape matching, surface descriptor

1. Introduction

The comparison of shapes is a crucial task in shape analysis. E.g. regarding shape collections we want to be able to compare shapes from the same class to each other. For symmetry detection within objects we need to find adequate matches between points which re-

quire an appropriate point signature. Also, tasks like classification or shape segmentation require signatures which capture local surface properties as well as global shape information. These are just a few examples among many.

A vast amount of local surface signatures and global descriptors

have been introduced. In recent shape analysis approaches these descriptors are often combined (e.g. [KHS10]) by learning class specific weights for each descriptor in tasks like segmentation or classification. Although the accuracy of learned approaches is quite high they always require pre-labeled data for the training of the weights. Furthermore, the more descriptors need to be trained the more pre-labeled data has to be provided for accurate results. This underlines the relevance of accurate geometric descriptors.

Recently, Carrière et al. showed that augmenting local descriptors with topological information can increase the performance for segmentation and shape labeling tasks [COO15]. Still the topological information alone is not sufficient to distinguish local surface points. Hence, a descriptor which incorporates both local and topological surface information is a desirable goal.

In this paper we built upon the work of Tevs et al. [TBW*11] who grow geodesic discs around a surface point p in a small geodesic radius (5% of the longest bounding box edge of a shape) and measure the length of the isocontours at different radii. This captures local surface properties. To also describe the global shape of the object and capture major topological changes we exploit the global representation presented by Sun and Abidi [SA01], who compute the geodesic distances on the entire shape. Accordingly we measure the lengths of isocontours across the entire shape. When isocontours merge or split we are able to capture major topological events with this signature. Furthermore, we show that varying the sampling density of the selected radii vastly improves matching performance of this descriptor. Our contribution can be summarized as follows:

- We generalize the idea of Tevs et al. [TBW*11] and Sun and Abidi [SA01] by presenting the geodesic iso-curve signature (GICS), a surface descriptor with an intuitive geometric understanding, and proposing an appropriate sampling and distance function, such that the descriptor can be used for surface matching effectively. Furthermore, this signature can be used to describe sets of points, polygonal curves, surfaces or arbitrary combinations of the named geometric primitives.
- We compare the GICS to the wave kernel signature [ASC11], the Histogram of Geodesics descriptor presented in [LGB*13], and the local isocurve descriptor published by Tevs et al. [TBW*11] in several tests on shape collections which are provided with ground-truth correspondences and receive superior results regarding the correspondence quality.
- We show that the proposed descriptor is more robust to decimation than previous methods, which allows computations at a lower resolution.
- Furthermore, we present two applications. First, since we can extend the formulation of the GICS to arbitrary sets of geometric primitives, we are able to define signatures for line features. Secondly, we show that the GICS can be applied for value measured maps by using it as a comparability function in the soft maps framework presented in [SdGP*15].

2. Related Work

Over the last decades a lot of surface signatures have been presented. It is out of the scope of this paper to discuss all of these

descriptors, hence we will only mention some of the important representatives.

Local Descriptors Early approaches present descriptors which are invariant under rigid motion. Some examples are Shape Context [BMP00], Spin Images [JH99], integral volume descriptors [MHYS04, GMGP05], or multi-scale local features [LG05]. Mortara et al. [MPS*03] present a descriptor where they grow Euclidean balls around a surface point and measure the length of the surface intersection for different radii. Although this signature is not invariant to isometric transformations, the idea of this signature is similar to the one presented here.

Furthermore, a set of view based descriptors have been presented. Sun et al. project geodesic iso-curves onto the tangent plane at a surface point p and compare these fingerprints by computing the normal and projected contour length deviation along the contours [SA01, SPK*03]. As an intermediate representation they present a data structure, the global fingerprint, that stores the geodesic distance of the entire shape for a point p . In this work we will also sample the radii from this global representation, but in contrast to Sun et al. we will not project the iso-curves but directly compare the arclengths. Also, Mahiddine et al. present a similar idea by comparing level curves, which are 2D planar curves that are projected onto cutting planes [MMDmB14]. A further multi-resolution view based approach was presented by Dinh et al. [DK06]. They define multi-resolution Spin-Images. A disadvantage of projecting surfaces or curves onto 2D planes is that several 3D points might be projected onto the same 2D point, which can limit the descriptiveness. This is why we measure the lengths of the isocurves in 3D.

Several approaches were then proposed that are invariant under isometric transformations. Since geodesic distances are isometric invariant, a set of approaches based on geodesic distances has been proposed. E.g. Hilgaga et al. [HSKK01] integrate the geodesic distance from a surface point p to all other points. Elad et al. [EK03] apply a multidimensional scaling procedure to geodesic distances. The histogram of geodesic distances descriptor (HOG) [MS09, LGB*13] counts the number of sample points on the surface which fall in the same bin of geodesic distances measured from a local surface point. The separations between the bins are evenly spaced and a χ^2 -distance is used to compute the distance between two histograms. Compared to the descriptor we present here a surface integral is approximated between different geodesic radii. Hence, a lot of sample-points are required to obtain a good approximation. Gatzke et al. [GGGZ05] compute curvature values within evenly spaced geodesic rings. Similarly, Tevs et al. [TBW*11] present a descriptor which computes the lengths of geodesic isocurves in a small local radius around a surface point p . We extend their work by computing geodesic radii on the entire shape and show that by using an adequate sampling function of the radii, this descriptor can outperform current local surface descriptors with respect to their matching performance.

Furthermore, spectral descriptors like the heat-kernel signature (HKS) or wave-kernel signature (WKS) have been presented which are based on the Eigen decomposition of the Laplace-Beltrami matrix of a 3D object [SOG09, BK10, ASC11]. These spectral descriptors are appealing since they characterize a surface point in rela-

tion to the entire shape. Especially, for the WKS the amount of frequencies can be preselected which allows more accurate matching than previously proposed methods. Our geodesic iso-curve signature also puts the local surface points in context to the global shape. Especially, since prominent topological events are captured in the signature more accurate matching results are achieved for the test repositories. Furthermore, in contrast to spectral descriptors our signature can be extended to characterize not only points but also polygonal curves, surfaces, as well as combined sets of these primitives. We will also show that our descriptor is more robust under decimation.

Finally, methods which are based on well known descriptors from image analysis have been presented. E.g. Kokkinos et al. [KBLB12] presented a shape context descriptor which sorts descriptors like the WKS or HKS into different spherical and radial bins around a surface point.

Topological Descriptor Carriere et al. published a topological signature for points on 3D shapes [COO15]. They grow geodesic balls around a surface point p and intersect this ball with the surface. The radii for which the number of holes in the intersected surface changes are stored in a persistence diagram, allowing to capture topological object features. Carriere et al. show that augmenting local descriptors with this topological signature improves matching and supervised shape labeling results. This is the motivation for our approach. We present a descriptor which by design captures both local surface properties and topological events.

Supervised Learning Recently, several approaches based on supervised learning have been presented. For this, a set of positive and negative point correspondences are defined on training shapes. These prelabelled correspondences are input to a neural network, which learns required parameters to map points on the mesh to the descriptor space (see e.g. [BMM*15, BMR*16]). These methods allow to learn very precise class specific descriptors but always require prelabelled data from objects of the same shape class.

Correspondence by Measure Valued Maps A recent line of research focuses on the relaxation of the point-to-point correspondence problem, which computes a map $m : \mathcal{M} \rightarrow \mathcal{M}'$ for two shapes \mathcal{M} and \mathcal{M}' to the generation of measure valued maps $m : \mathcal{M} \rightarrow \text{Prob}(\mathcal{M}')$ (see e.g. [SdGP*15, SNB*12, OBCS*12]). The measure valued maps are usually based on a point compatibility function like the WKS. Hence, our descriptor could be plugged into such a framework to generate measure valued maps. As an application we insert our descriptor as a compatibility function into the softmaps framework [SdGP*15].

Global Signatures Further research focuses on the definition of global signatures, which allow to compare entire shapes. E.g. Reuter et al. define the shape-DNA which stores the eigenvalues of the Laplace operator [RWP06]. Rustanov et al. compare histograms of pairwise distances in the descriptor space of the global point signature [Rus07]. These are just a few examples among many.

3. Properties of the Geodesic Iso-Curve Signature

Overall, the GICS descriptor has the following properties:

- **Invariance to isometric transformations (intrinsic):** The computation of the lengths of the iso-curves is based on the geodesic distances from a surface point x . The geodesic distance $g(x, y)$ on points is an intrinsic metric (cf. [MPWC13]), i.e. for an isometry $T : \mathcal{M} \rightarrow \mathcal{M}'$ the identity $g(x, y) = g(T(x), T(y))$ holds. Then for all points y on an isocontour $c_x(r_i)$ (with length $l_x(r_i)$) of the source point x at radius r_i it holds that $g(x, y) = g(T(x), T(y))$. Furthermore $g(y_i, y_j) = g(T(y_i), T(y_j))$ for all $y_i, y_j \in c_x(r_i)$. Hence, we measure the same length $l_{T(x)}(r_i)$ of an isocontour $c_{T(x)}(r_i)$ for the transformed point $T(x)$ at radius r_i , i.e. $l_{T(x)}(r_i) = l_x(r_i)$ (More details on the computation of the radii are given in Section 4).
- **Sensitive to topological and geometric features:** Upon topological events geodesic isocontours split, merge, or parts of the isocontour vanish (when a geometric feature ends). When the isocontours that grow smoothly around a surface point meet, they are merged, or respectively split when they grow around separate topological features. In many cases such topological events cause a sign change in the gradient of the signature function such that they can be detected as peaks (see Figure 2). Furthermore, when the radius exceeds the length of a geometrical feature (e.g. the fingers of the humans in Figure 1), the length of the radius decreases and drops to zero for the respective geometric part. Hence, these events induce local maxima in the local continuous length-function of the isocontours for prominent topological features. This way characteristic topological changes of a shape are captured.
- **Capture local surface properties:** Local samples of the radii around a surface point capture the local shape properties. By adjusting the distribution of samples (geodesic radii) in the domain of the length function $l_x(r_i)$, we can make sure that the local shape properties are captured well.
- **Stability under re-tessellation and mesh decimation:** The geodesic descriptor is stable under decimation of the mesh. In Section 5 we show that the signature distance of sub-sampled versions of the mesh compared to the original shape are smaller than for comparable signatures.
- **Descriptor for arbitrary subsets of a shape:** The GICS can be used to describe arbitrary sets of points, polygonal curves, polygonal graphs, or a combination of these. In Section 6.1 we show an application for matching line features.
- **Intuitive geometric understanding**
- **Easy implementation:** The main part of the implementation involves the computation of the geodesic distances, while the implementation of the descriptor itself is extremely simple.

4. The Geodesic Iso-Curve Signature

Given a manifold surface \mathcal{M} , our geodesic iso-curve signature is based on the computation of lengths $l_x(r_i)$ of isocontours around a surface point x on \mathcal{M} , at n different geodesic radii r_i :

$$D_x = [l_x(r_1), l_x(r_2), \dots, l_x(r_n)] \in \mathbb{R}^n$$

with radii

$$r_1 \leq r_2 \leq \dots \leq r_n, r_i \in \mathbb{R}.$$

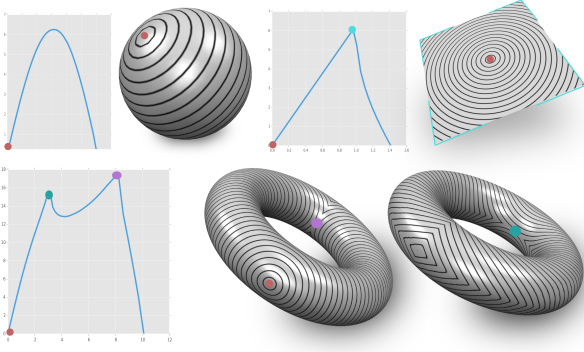


Figure 2: Signatures of a sphere, plane and a torus: The red dots represent the source points x from which geodesic disks are grown. On the sphere no topological events occur. The cyan dot on the signature of the plane denotes the point at which the radii grow out of the plane, and parts of the isocontours vanish. The green point on the torus marks the first topological event, where the isocontour is split. The purple dot is positioned at the second topological event. Here, the two isocontours merge again. On the left the signature with the respective events are depicted. The two images of the torus show the evolution of the isocontours on the shape from the front (left) and the back (right).

The function $l_x(r_i)$ maps the radius r_i to the length of the corresponding iso-curve:

$$l_{x \in \mathcal{M}} : [0, d_{\max}(x)] \rightarrow \mathbb{R},$$

where $d_{\max}(x) = \max\{g(x, y) | \forall y \text{ on } \mathcal{M}\}$ denotes the maximal geodesic distance on \mathcal{M} from a surface point x .

For the computation of the signature several design choices have to be made: (1) an algorithm to compute the geodesic distances is selected, and (2) a method for the approximation of the isocontours of the geodesic discs has to be chosen. In addition (3), a distance metric has to be defined on the signatures for comparison. Finally, (4) we need to decide how to sample the geodesic radii on \mathcal{M} . In the following we will discuss these design choices.

4.1. Geodesic Distances

Diverse algorithms exist to compute geodesic distances on a triangular surface mesh \mathcal{M} (e.g. [CWW13, CHK13, BK07, SSK*05]). We use approximate exact geodesic distances [BK07], because of their robustness. However, if performance is the main goal the fast heat method [CWW13] would be a good choice and depending on the type of input data there might other preferable options. E.g. we can apply anisotropic distance fields [CHK13] if we want to capture different orientations as in [BMR*16].

4.2. Radii Computation

We compute the geodesic distance from a seed point x on \mathcal{M} . To approximate the length $l_x(r_i)$ of an isocontour corresponding to a radius r_i we sum up the lengths of the line segments between the two intersection points of each triangle intersecting the isocontour

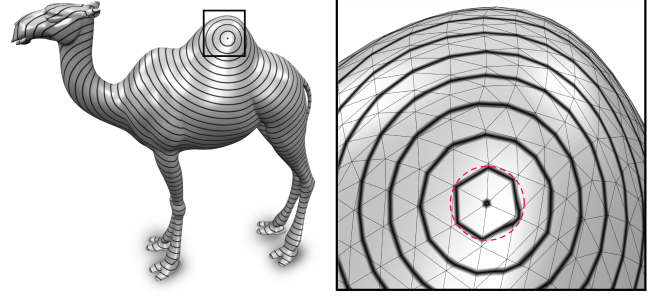


Figure 3: Approximation of the length of an isocontour (red). The sub arcs of the isocontour are approximated by the line which connects the two points at which the isocontour intersects the triangle of a mesh \mathcal{M} (left).

(see Figure 3). Let $T = (v_0, v_1, v_2)$ denote a triangle that is intersected by the isocontour along the edges (v_0, v_1) and (v_0, v_2) for a radius r_i . The segment $l_{x_T}(r_i)$ of the isocontour that intersects this triangle is approximated as:

$$l_{x_T}(r_i) = |p_1 - p_2|$$

with

$$p_j = (1 - \alpha_j) \cdot v_0 + \alpha_j \cdot v_j \quad \text{and} \quad \alpha_j = \frac{|r_i - g_0|}{|g_j - g_0|},$$

where g_i denotes the geodesic distance between x and v_i . The sum of the lengths of the segments $l_{x_T}(r_i)$ approximates the length of the respective isocontour.

4.3. Distance Computation

For two shapes \mathcal{M} and \mathcal{M}' , we compute the distance of the signatures of two points x on \mathcal{M} and y on \mathcal{M}' . Similar to [ASC11], we define a distance using the L^1 norm of normalized iso-curve length distances at the same radii to compare the geodesic signatures:

$$D_g(x, y) = \sum_{i=0}^{i_{\max}} \frac{|l_x(r_i) - l'_y(r_i)|}{|l_x(r_i) + l'_y(r_i)|}$$

The lengths $l'_y(r_i)$ on the shape \mathcal{M}' are interpolated linearly from the the two radii r'_j and r'_{j+1} with $r'_j \leq r_i \leq r'_{j+1}$, since the radii are not necessarily sampled at the same values (cf. Section 4.4). From each point the maximal geodesic distance $d_{\max}(x)$ on \mathcal{M} can be different. Hence, we compute a maximal index i_{\max} for which both signatures have defined values. Hence i_{\max} is defined as

$$i_{\max} = \max\{i | r_i \leq \min(d_{\max}(x), d_{\max}(y))\}$$

Finally, to make this distance symmetric, we compute the descriptor distance $D_{GICS}(x, y)$ as:

$$D_{GICS}(x, y) = \frac{D_g(x, y) + D_g(y, x)}{2}.$$

4.4. Radius Sampling

The function $l_{x \in \mathcal{M}} : [0, d_{\max}(x)] \rightarrow \mathbb{R}$, which describes the length of an isocontour at radius r needs to be sampled at discrete values r_1, \dots, r_n in order to obtain a descriptor of size n . An obvious choice




Figure 4: Radius sampling. To capture local shape properties and guarantee local shape awareness we adjust the distribution of samples (geodesic radii) in the domain of the length function according to a monotonically increasing sampling function (right). On the left we depict the respective sample distribution on a plane (red).

would be to uniformly sample r between 0 and $d_{\max}(x)$. Unfortunately, then the local shape is not captured well by the signature. To achieve local shape awareness we need to increase the density of local samples. Hence, we introduce a continuous monotonically increasing sampling function $f : [0, l_s] \rightarrow \mathbb{R}$. The domain of this sampling function is fixed to preset values $[0, l_s]$ on which we uniformly sample (see x -axis in Figure 4 for $n = 4$). This domain is equal for each shape, since we want to apply the same sampling for each signature. Hence, the following increment is computed in the domain:

$$\delta = \frac{l_s}{n}$$

Since the maximum function value is $f(l_s)$, in order to obtain radius values the function values need to be rescaled by $\frac{d_{\max}(x)}{f(l_s)}$, such that the radii can be computed as:

$$r_i = f(i \cdot \delta) \cdot \frac{d_{\max}(x)}{f(l_s)}$$

In our experiments we use sampling functions of $f(x) = x$, $f(x) = x^2$, and $f(x) = x^4$. Overall, we found that applying the sampling function $f(x) = x^4$ allows a more accurate distinction between local surface properties.

4.5. Descriptor for Arbitrary Subsets

[BK07] present an extension to the exact geodesics computation in [SSK*05]. While the latter compute geodesic distances from isolated point sets, Bommes et al. extended their work for arbitrary, possibly open, polygons on the mesh which define the zero set of the distance field. Hence, arbitrary sources can be set for the computation of geodesic distances. E.g. we can choose point sets, curves, polygonal graphs or even a mix of the afore mentioned as sources. Then, the isocontours of the geodesic radii can be computed for these sets as well, so that we can generate signatures for sets of primitives:

$$D_{\Omega} = [l_{\Omega}(r_1), l_{\Omega}(r_2), \dots, l_{\Omega}(r_n)],$$

where Ω denotes the set of primitives which is defined to be the zero set of the distance field. The geodesic iso-curve signature of Ω can then be computed exactly in the same manner as described above.

5. Evaluation

We evaluate the geodesic signature by comparing it to the previously proposed wave kernel signature (WKS) [ASC11], the histogram of geodesic distances (HOG) [LGB*13], and the original signature presented by Tevs et al. [TBW*11]. We chose the WKS, since it is considered a state-of-the-art geometric descriptor and is incorporated in many recent approaches (e.g. [JCK15, LB14, SNB*12]). Litman et al. state that the WKS achieves state-of-the-art performance in many deformable shape analysis tasks [LB14]. The HOG signature and the descriptor by Tevs et al. are those most related to our work. We can measure the matching performance of these descriptors by comparing the correspondence scores on shape collections with pre-labeled ground-truth data. Furthermore, we present a qualitative evaluation and show that our descriptor is more robust under decimation.

5.1. Test Setup

Following [BMR*16] we compute the cumulative match characteristic (CMC), the receiver operator characteristic (ROC) and the Princeton test on the shape collections FAUST [BRLB14], SCAPE [ASK*05], and TOSCA [BBK08]. These shape collections are available online and have pre-labeled correspondences. Hence, we are able to perform a qualitative evaluation of our method.

Given a shape collection with a reference shape \mathcal{M}_0 with vertices V_0 and a set of shapes $\mathcal{M}_1, \dots, \mathcal{M}_m$ with vertex sets V_1, \dots, V_m respectively, which have pre-labeled ground-truth correspondences (v_{0_x}, v_{i_x}) for $v_{0_x} \in V_0$ and $v_{i_x} \in V_i, i \in \{1, \dots, m\}$ we perform the following tests:

- **CMC:** The CMC measures the probability of a correct correspondence among the k nearest neighbors, i.e. it captures the hit-rate of positive matches in the k -NN neighborhood in the descriptor space. To calculate the CMC we compute the $k = 100$ nearest neighbors for each vertex $v \in V_0$ on all shapes $\mathcal{M}_1, \dots, \mathcal{M}_m$. In case the correct correspondence is the i -th nearest neighbor with $i \leq k$, it is added to the characteristic at indices $\geq i$.
- **ROC:** The ROC measures the percentage of positive and negative matches. I.e. for positive $c^+ = (p^+ \in V_0, q^+ \in V_i) \in C^+$ and negative matches $c^- = (p^- \in V_0, q^- \in V_i) \in C^-$ the false positive and true positive rates are measured as:

$$N_{tp} = \frac{|\{(p^+, q^+) \in C^+ | d(p^+, q^+) \leq \tau\}|}{|C^+|}$$

$$N_{fp} = \frac{|\{(p^-, q^-) \in C^- | d(p^-, q^-) \leq \tau\}|}{|C^-|},$$

where $d(p, q)$ measures the descriptor distance of the points p and q . The value τ defines a distance threshold in the descriptor space, which is increased from zero to the maximal descriptor distance.




Figure 5: Distance maps in descriptor space. A point on the reference shape (leftmost) is compared to all other points on the same and on other shapes. The shapes are taken from the FAUST dataset. Source points are located on the back of the hand, on the tip of the nose, and on the foot. For better visualization the values are cropped at the median.

- **Princeton protocol [KLF11]:** The Princeton protocol computes the percentage of matches that are within a geodesic radius r from the ground-truth correspondence. The geodesic distances are normalized by the square root area of the shape.

Furthermore, we use the following shape collections for our tests:

- **FAUST [BRLB14]:** This data set consist of ten humans which were scanned in ten different poses each. We use shape 80 as a reference shape to meshes 81 to 99. We perform the same test with our GICS and with the WKS.
- **TOSCA (from SHREC 2010, [BBK08]):** This shape collection consists of 64 objects, including 11 cats, 9 dogs, 3 wolves, 8 horses, 6 centaurs, and two different male figures, containing 7 and 20 poses. We down-sampled the meshes from 27.000 to about 5.000 vertices. We used the first shape of each category as the reference shape.
- **SCAPE [ASK*05]:** The SCAPE data set consists of human scans in seventy different poses. We use the first fifty meshes for our tests.

Furthermore we use the following parameters for our tests. Note

that we use the descriptors with the parameters as they are presented in the respective paper:

- **GICS:** We use a descriptor size of 100, i.e. we take 100 samples of geodesic radii. The sampling functions applied here are $f(x) = x$, $f(x) = x^2$, and $f(x) = x^4$. The value l_s is set to 2.
- **WKS:** As described in [ASC11] we compute a descriptor size of 100 with 300 eigenvalues of the Laplace Beltrami matrix.
- **HOG:** We compute 100 bins with equal geodesic spacing. The number of samples is in the size of the mesh (i.e. number of vertices).
- **Tevs et al.:** We compute 16 geodesic radii, with the maximal geodesic distance equal to 5% of the longest bounding box edge. We also performed tests with more samples (100 radii), which gave poorer results.

5.2. Results

Figure 5 shows a qualitative evaluation of our signature. We visualize the descriptor distance exemplary for meshes from the FAUST dataset. The signature distances are computed across the shapes for

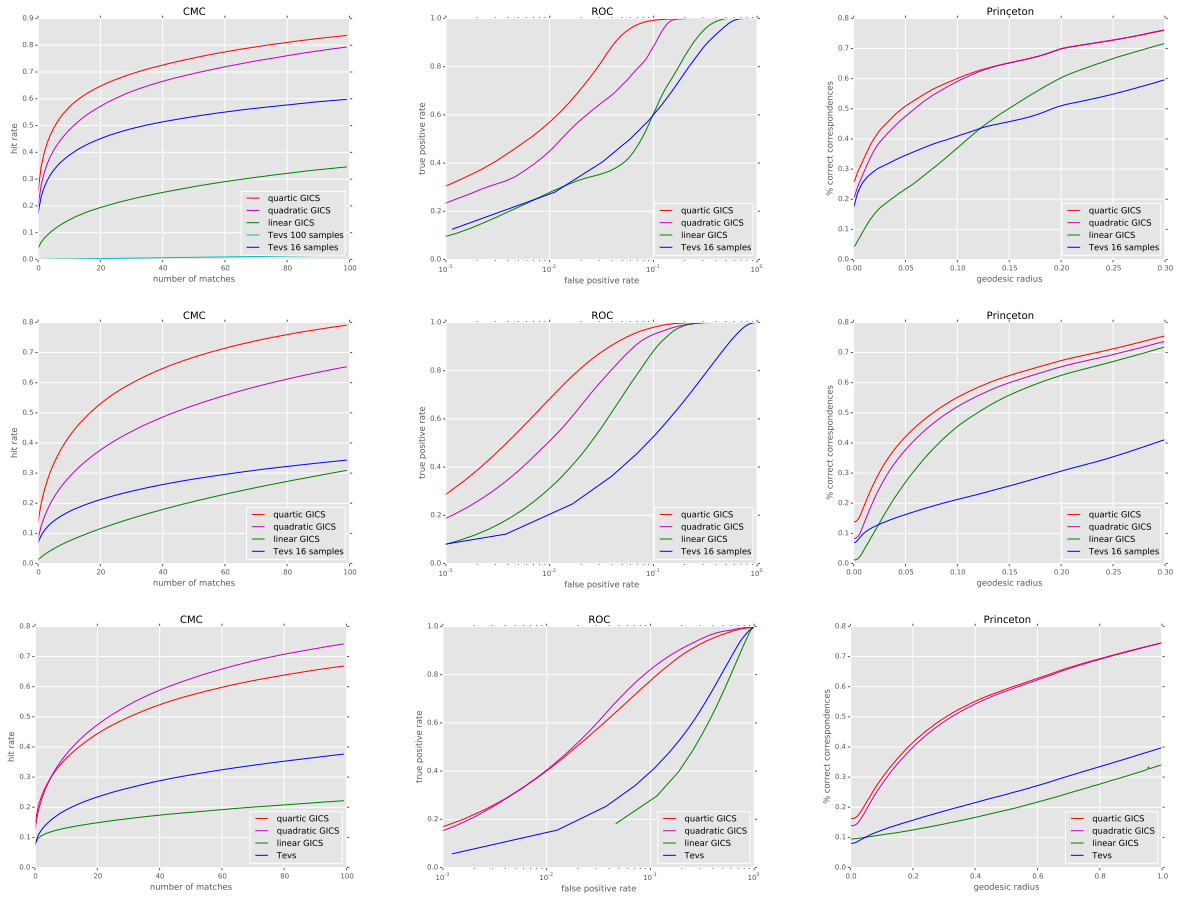


Figure 6: Results of the CMC, ROC, and Princeton test (from left to right) on the dataset FAUST, SCAPE, and TOSCA (top to bottom) for the signatures GICS (with sampling function $f(x) = x$, $f(x) = x^2$, and $f(x) = x^4$) and by Tevs et al.

a point selected on the hand (top), the nose (middle), and on the foot (bottom).

In Figure 6 we evaluate different sampling functions and compare the matching results to the signature presented by Tevs et al. For this, we show the plots for the CMC for $k = 100$, ROC, and Princeton test with a maximal geodesic radius of 0.3 for the shape collections FAUST, SCAPE and TOSCA (top to bottom). We observe that we achieve much more accurate matching performance for the sampling functions $f(x) = x^2$ and $f(x) = x^4$ compared to uniformly spaced sample radii (Tevs and linear GICS).

Figure 7 depicts the respective plots for the GICS (sampling $f(x) = x^4$), the WKS and the HOG. We perform the tests both symmetrically (symmetric counter parts of a correspondence are identified as a correct match) and asymmetrically. We compute the symmetric counter parts of each vertex by defining a symmetry plane on the source shapes and aligning the two halves of the mesh via the iterative closest points method. Note, that the initial shape of the FAUST and the SCAPE dataset have slight isometric deformations, hence the symmetric matches are not as accurate as for the TOSCA dataset. All symmetric results are depicted as dashed curves. We can observe that our geodesic iso-curve descriptor out-

performs the previous signatures in nearly all tests. Especially, for the CMC significant improvements can be measured. Furthermore, our matching results can also compete with current approaches in supervised learning (cf. [BMR*16]). Although our results are not as accurate as theirs, we do not require a training set, while still receiving comparable results.

5.3. Stability under Re-Tessellation and Mesh Decimation

We demonstrate that the GICS is more stable under mesh decimation compared to the wave kernel signature. Figure 8 depicts three meshes with their respective descriptor distance to different lower resolution versions of the meshes. The descriptor distances from each point to the closest point on the high resolution shapes are computed. We compute normalized signature distances for both our geodesic iso-curve signature (top row) and the wave kernel signature (bottom row). The white color indicates a distance of zero (i.e. zero error), while large distances are marked red (large error). The first two meshes are taken from the FAUST data set. To the left the original scan data with 176.744 vertices is depicted. It is compared to its respective registration mesh with 6980 vertices (first column) and a low resolution version with 500 vertices. The distance values

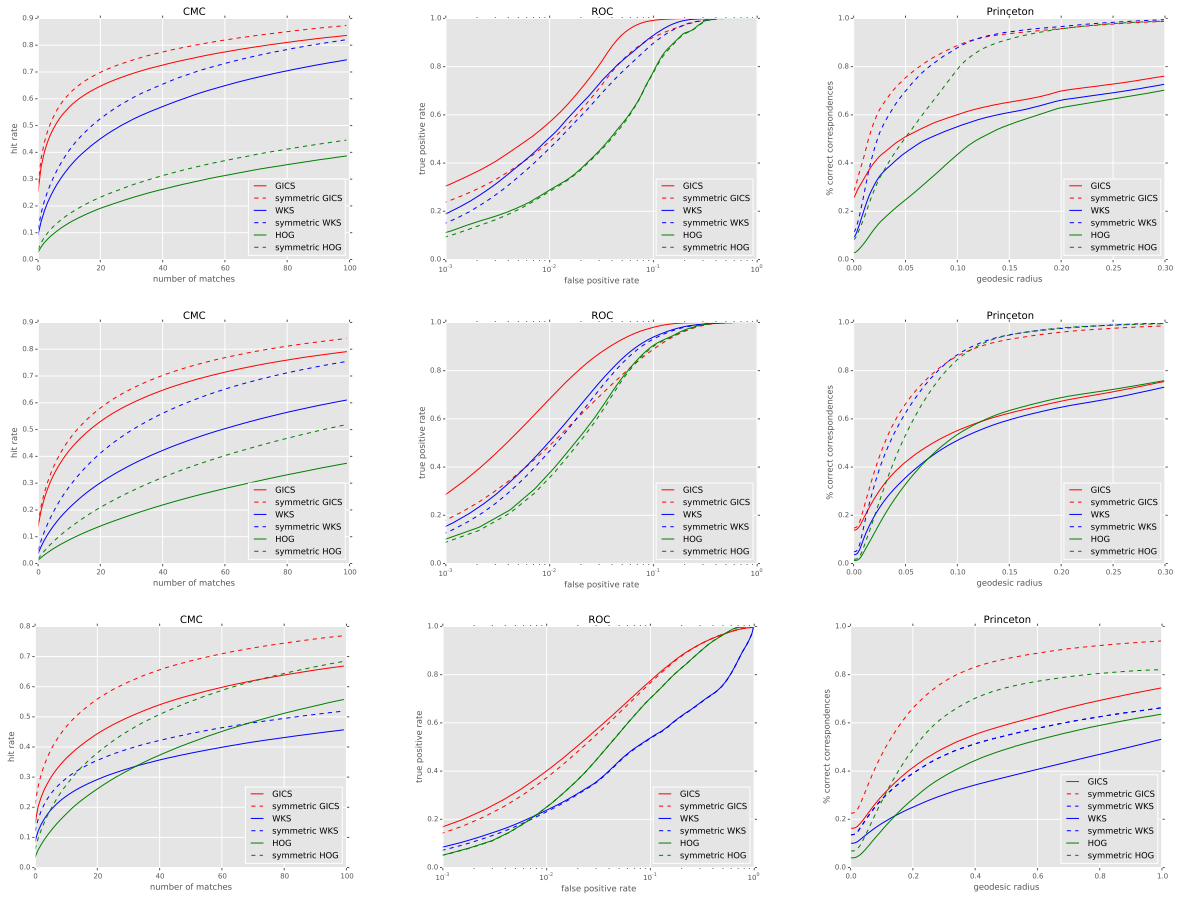


Figure 7: Results of the CMC, ROC, and Princeton test (from left to right) on the dataset FAUST, SCAPE, and TOSCA (top to bottom) for the descriptors GICS (sampling function $f(x) = x^4$), WKS, and HOG. The dashed curves describe symmetric results (i.e. the symmetric counter parts of a correspondence are considered as correct matches) while the solid curves depict the asymmetric results. We can observe that our descriptor outperforms the WKS and the HOG in nearly all tests.

for our geodesic iso-curve signature range from 0 to 0.47, while the distance values for the WKS range from 0.85 to 1. In the images above we show the complete range (0 – 0.47) for our signature, while we depict a range of 0.95 – 1 for the WKS. We also compare the FAUST mesh with 6980 vertices to two lower resolution versions (1000 and 500 vertices) with values ranging from 0 to 0.14 for the GICS and 0 to 0.29 for the WKS. Finally, we also show two lower resolution versions of the elephant with values in 0 to 0.26 for the GICS and 0 to 0.39 for the WKS. Overall, we observe that the wave kernel signature shows a greater error, i.e. the maximum distance is approximately twice as large as the error for the geodesic iso-curve signature. Furthermore, the error values of the GICS seem to be quite constant on the meshes, whereas those of the WKS show several local maxima on the shapes.

6. Applications

In the following we will present two applications for our geodesic iso-curve descriptor.

6.1. Feature Line Signature

Arbitrary sources can be set for the computation of geodesic distances. This is especially useful for the description of curves, which allows us to define a signature for line features.

In the case of line features we have one signature for each feature curve. Figure 9 illustrates an example of the geodesic signature on feature lines for a "tele alien" shape. The initial feature lines are computed with [YBS05] (left). We computed the distance of the red feature on the second tele alien to all other line features of the mesh and show its best matching feature curve on the third shape. The bottom row encodes the distance values where we cutoff the distance threshold at increasing values (left to right).

Further, we show an example for across shape feature line matching in Figure 10. Here the circular shaped feature around the left eye of the first tele alien is selected as the source feature. The distance values to the other features are encoded in red (low distance) and blue (high distance). The best matches on second shape also represent circular arcs around the tele aliens eyes.




Figure 8: Stability under decimation test. Each of the three meshes is downsampled to two lower resolutions. The descriptor distances from each point to the closest point on the high resolution shapes are computed. High distance values are encoded in red, while points with low distance values are colored white. I.e. strong error is marked red. The upper row shows the distance values for our descriptor, while the bottom row shows distances for the WKS. The maximal distances of the WKS are about twice as high.

6.2. Softmaps

Solomon et al. have presented an efficient method to compute soft maps between surfaces in [SdGP*15]. These soft maps describe a measure valued map $m : \mathcal{M} \rightarrow \text{Prob}(\mathcal{M}')$, where each point on \mathcal{M} is mapped to a probability distribution over \mathcal{M}' . They are based on a compatibility function $c(x, y)$, where small values of $c(x, y)$ indicate that x on shape \mathcal{M} and y on shape \mathcal{M}' are geometrically similar. We use the distance function described in Section 4, which computes descriptor distances on our signatures as a compatibility function. Figure 11 shows soft maps from the leftmost shape to itself (left) and to two other shapes for two different points x on the source shape (on the finger and the knee).

We compare the soft maps generated with our geodesic iso-curve signature to soft maps applying the WKS as a compatibility function in Figure 12. The selected source point (finger) equals the one chosen in Figure 11 (top). The probability distributions computed across the shapes differ more than those computed with our signature (especially for the female shape). Furthermore, we observe strong maxima at the feet of the third shape. Overall, we believe that the distributions computed with our signature enable good results in the soft maps framework and generate similar distributions across shapes from the same collection.

Limitations The presented descriptor has some limitations which need to be pointed out. First of all, to compute the signature we

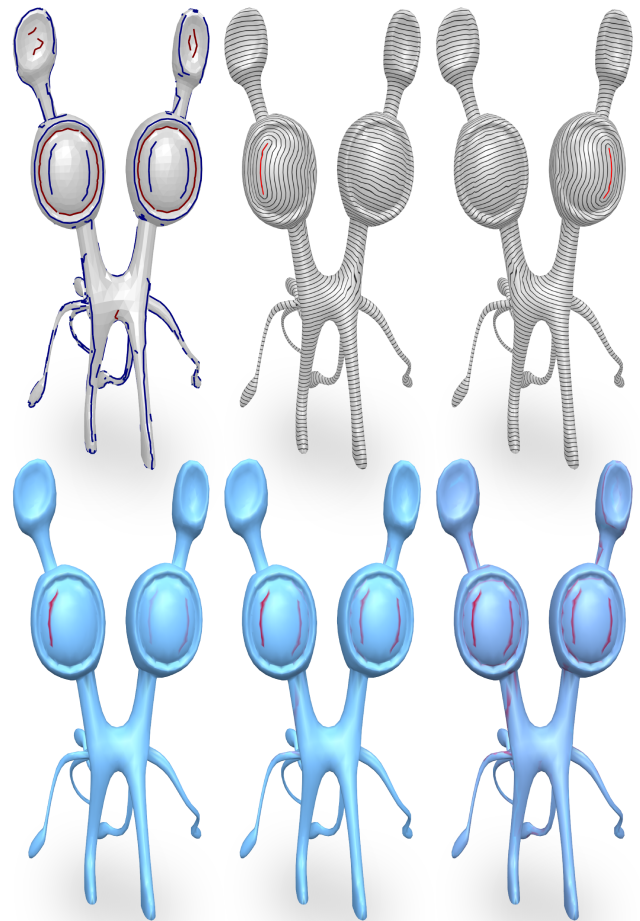


Figure 9: Extension to line features. Geodesic isocontours are computed for radii growing around a curve segment. We computed feature curves with [YBS05] (top, left) and illustrate the isocontours around a representative curve (top, middle) and its best match in descriptor space (top, right). The bottom row encodes the distance values of the feature curves in descriptor space, where the maximal descriptor distance is cut off for increasing values (left to right). The red color indicates a high matching score (i.e. low descriptor distance).

need to compute geodesic distances from all sample points. This affects the computational efficiency. E.g. to compute the signature with exact geodesics [BK07] for the FAUST datasets (6890 vertices) takes about 43 ms per sample. We can improve this by using geodesics in heat [CWW13] (about 3 ms). Also, it might be interesting to evaluate whether Biharmonic distances presented in [SRGB14] can be used to approximate the geodesic distances in this setting to improve the performance.

A further limitation of our approach is that this signature is not scale invariant in its current form, we believe that it can be transformed into a scale invariant version by scaling the meshes by an average geodesic distance before the computation. We leave this task for future work.

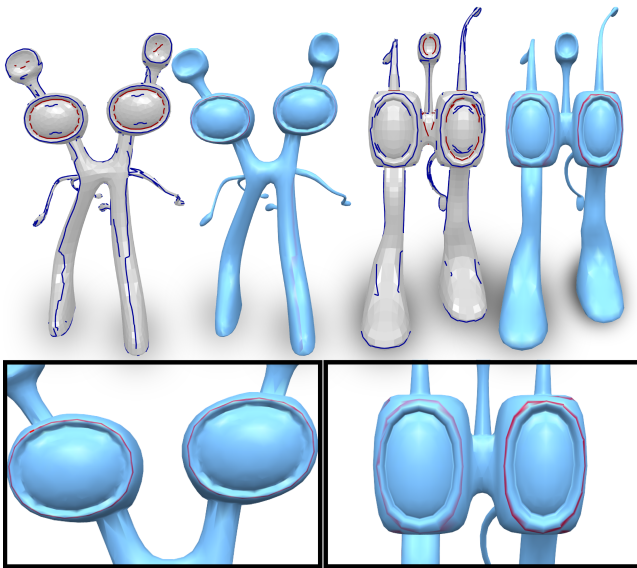


Figure 10: Across shape feature line matching. The circular shaped feature curve around the left eye of the first tele alien is selected as the source feature. The distance values to the other features are encoded in red (low distance) and blue (high distance). The best matches on second shape represent circular arcs around the tele aliens eyes as well.

7. Conclusion

In this paper we have extended the signatures presented in [TBW*11] and [SA01] to capture both prominent topological information and local surface properties of a 3D object. The geodesic iso-curve signature allows to describe arbitrary sets of geometric primitives, which can lead to interesting new research directions in the field of shape analysis. We have performed extensive tests on shape repositories, which result in superior matching scores compared to the WKS, HOG, and the signature by Tevs et al. Also, we have shown that our descriptor is more robust to decimation of the mesh than the WKS.

Furthermore, we have demonstrated two applications for our point signature. We showed that it can be applied for the description of feature lines and that it enables good results in the framework of soft maps.

Especially, due to the significant results in the matching tests, we believe that this descriptor can have important impact on shape analysis tasks. Not only does it describe local surface points more accurately than currently published local descriptors, also in the field of supervised learning a more descriptive signature can reduce the amount of training data that is required for the respective task.

Acknowledgements

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement n^o [340884]. Several models are provided courtesy of INRIA by the




Figure 11: Soft maps. We use our descriptor as a compatibility function for the soft maps framework presented in [SdGP*15]. The soft maps are computed from the source shape (left most) to itself and to two other shapes. The reference points for the depicted maps are selected on the finger and on the knee of the source object.

AIM@SHAPE-VISIONAIR Shape Repository. We would like to thank Jan Möbius for creating and maintaining the geometry processing framework OpenFlipper [?] as well as the reviewers for their insightful comments.

References

- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on* (Nov 2011), pp. 1626–1633. doi:10.1109/ICCVW.2011.6130444. 2, 4, 5, 6
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: Scape: Shape completion and animation of people. *ACM Trans. Graph.* 24, 3 (July 2005), 408–416. doi:10.1145/1073204.1073207. 5, 6
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*, 1 ed. Springer Publishing Company, Incorporated, 2008. 5, 6
- [BK07] BOMMES D., KOBELT L.: Accurate computation of geodesic distance fields for polygonal curves on triangle meshes. *VMV 7* (2007), 151–160. 4, 5, 9




Figure 12: Comparison of soft maps to the WKS. We use the WKS as a compatibility function for the soft maps framework presented in [SdGP*15]. The selected source point (finger) equals the one chose in Figure 11(top). We can observe that probability distributions computed across the shapes differ more than those computed with our signature (especially for the female shape). Furthermore we observe strong maxima at the foot of the third shape.

- [BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on* (June 2010), pp. 1704–1711. doi:10.1109/CVPR.2010.5539838. 2
- [BMM*15] BOSCAINI D., MASCI J., MELZI S., BRONSTEIN M. M., CASTELLANI U., VANDERGHEYNST P.: Learning Class-specific Descriptors for Deformable Shapes Using Localized Spectral Convolutional Networks. *Computer Graphics Forum* (2015). doi:10.1111/cgf.12693. 3
- [BMP00] BELONGIE S., MALIK J., PUZICHA J.: Shape context: A new descriptor for shape matching and object recognition. In *NIPS* (2000), pp. 831–837. 2
- [BMR*16] BOSCAINI D., MASCI J., RODOLÀ E., BRONSTEIN M. M., CREMERS D.: Anisotropic diffusion descriptors. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 431–441. 3, 4, 5, 7
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (Piscataway, NJ, USA, June 2014), IEEE. 5, 6
- [CHK13] CAMPEN M., HEISTERMANN M., KOBBELT L.: Practical anisotropic geodesy. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2013), SGP '13, Eurographics Association, pp. 63–71. doi:10.1111/cgf.12173. 4
- [COO15] CARRIÁLRE M., OUDOT S. Y., OVSJANIKOV M.: Stable topological signatures for points on 3d shapes. *Computer Graphics Forum* 34, 5 (2015), 1–12. doi:10.1111/cgf.12692. 1, 2, 3
- [CWW13] CRANE K., WEISCHEDEL C., WARDETZKY M.: Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. Graph.* 32, 5 (Oct. 2013), 152:1–152:11. doi:10.1145/2516971.2516977. 4, 9
- [DK06] DINH H. Q., KROPAC S.: Multi-resolution spin-images. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)* (June 2006), vol. 1, pp. 863–870. doi:10.1109/CVPR.2006.197. 2
- [EK03] ELAD A., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 10 (Oct. 2003), 1285–1295. doi:10.1109/TPAMI.2003.1233902. 2
- [GGGZ05] GATZKE T., GRIMM C., GARLAND M., ZELINKA S.: Curvature maps for local shape comparison. In *International Conference on Shape Modeling and Applications 2005 (SMI'05)* (June 2005), pp. 244–253. doi:10.1109/SMI.2005.13. 2
- [GMGP05] GELFAND N., MITRA N. J., GUIBAS L. J., POTTMANN H.: Robust global registration. In *Proceedings of the Third Eurographics Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2005), SGP '05, Eurographics Association. URL: <http://dl.acm.org/citation.cfm?id=1281920.1281953>. 2
- [HKK01] HILAGA M., SHINAGAWA Y., KOHMURA T., KUNII T. L.: Topology matching for fully automatic similarity estimation of 3d shapes. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 203–212. doi:10.1145/383259.383282. 2
- [JCK15] JENI L. A., COHN J. F., KANADE T.: Dense 3d face alignment from 2d videos in real-time. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)* (2015). URL: articles/jeni15FG_ZFace.pdf. 5
- [JH99] JOHNSON A. E., HEBERT M.: Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Trans. Pattern Anal. Mach. Intell.* 21, 5 (May 1999), 433–449. doi:10.1109/34.765655. 2
- [KBLB12] KOKKINOS I., BRONSTEIN M. M., LITMAN R., BRONSTEIN A. M.: Intrinsic shape context descriptors for deformable shapes. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on* (June 2012), pp. 159–166. doi:10.1109/CVPR.2012.6247671. 3
- [KHS10] KALOGERAKIS E., HERTZMANN A., SINGH K.: Learning 3d mesh segmentation and labeling. In *ACM SIGGRAPH 2010 Papers* (New York, NY, USA, 2010), SIGGRAPH '10, ACM, pp. 102:1–102:12. doi:10.1145/1833349.1778839. 2
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. In *ACM SIGGRAPH 2011 Papers* (New York, NY, USA, 2011), SIGGRAPH '11, ACM, pp. 79:1–79:12. doi:10.1145/1964921.1964974. 6
- [LB14] LITMAN R., BRONSTEIN A. M.: Learning spectral descriptors for deformable shape correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 1 (Jan 2014), 171–180. doi:10.1109/TPAMI.2013.148. 5
- [LG05] LI X., GUSKOV I.: Multi-scale features for approximate alignment of point-based surfaces. In *Proceedings of the Third Eurographics Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2005), SGP '05, Eurographics Association. URL: <http://dl.acm.org/citation.cfm?id=1281920.1281955>. 2
- [LGB*13] LIAN Z., GODIL A., BUSTOS B., DAOUDI M., HERMANS J., KAWAMURA S., KURITA Y., LAVOÛAL G., NGUYEN H. V., OHBUCHI R., OHKITA Y., OHISHI Y., PORIKLI F., REUTER M., SIPHERAN I., SMEETS D., SUETENS P., TABIA H., VANDERMEULEN D.: A comparison of methods for non-rigid 3d shape retrieval. *Pattern Recognition* 46, 1 (2013), 449 – 461. doi:<http://dx.doi.org/10.1016/j.patcog.2012.07.014>. 2, 5
- [MHYS04] MANAY S., HONG B.-W., YEZZI A. J., SOATTO S.: *Computer Vision - ECCV 2004: 8th European Conference on Computer Vision, Prague, Czech Republic, May 11-14, 2004. Proceedings, Part IV*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, ch. Integral Invariant Signatures, pp. 87–99. URL: http://dx.doi.org/10.1007/978-3-540-24673-2_8. 2
- [MMDmB14] MAHIDDINE A., MERAD D., DRAP P., M. BOÑÁR J.: Partial 3d-object retrieval using level curves. In *Soft Computing and Pattern Recognition (SoCPaR), 2014 6th International Conference of* (Aug 2014), pp. 77–82. doi:10.1109/SOCPAR.2014.7007985. 2
- [MPS*03] MORTARA M., PATANé G., SPAGNUOLO M., FALCIDIENO B., ROSSIGNAC J.: Blowing bubbles for multi-scale analysis and decomposition of triangle meshes. *Algorithmica* 38, 1

- (Oct. 2003), 227–248. URL: <http://dx.doi.org/10.1007/s00453-003-1051-4>. 2
- [MPWC13] MITRA N. J., PAULY M., WAND M., CEYLAN D.: Symmetry in 3d geometry: Extraction and applications. *Computer Graphics Forum* 32, 6 (2013), 1–23. doi:10.1111/cgf.12010. 3
- [MS09] MAHMOUDI M., SAPIRO G.: Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models* 71, 1 (2009), 22–31. doi:<http://dx.doi.org/10.1016/j.gmod.2008.10.002>. 2
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4 (July 2012), 30:1–30:11. doi:10.1145/2185520.2185526. 3
- [Rus07] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the fifth Eurographics symposium on Geometry processing* (2007), Eurographics Association, pp. 225–233. 3
- [RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-beltrami spectra as shape-dna of surfaces and solids. *Computer-Aided Design* 38, 4 (2006), 342–366. 3
- [SA01] SUN Y., ABIDI M. A.: Surface matching by 3d point’s fingerprint. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on* (2001), vol. 2, pp. 263–269 vol.2. doi:10.1109/ICCV.2001.937634. 1, 2, 10
- [SdGP*15] SOLOMON J., DE GOES F., PEYRÉ G., CUTURI M., BUTSCHER A., NGUYEN A., DU T., GUIBAS L.: Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Trans. Graph.* 34, 4 (July 2015), 66:1–66:11. doi:10.1145/2766963. 2, 3, 9, 10, 11
- [SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft maps between surfaces. *Computer Graphics Forum* 31, 5 (2012), 1617–1626. doi:10.1111/j.1467-8659.2012.03167.x. 3, 5
- [SOG09] 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* (Aire-la-Ville, Switzerland, Switzerland, 2009), SGP ’09, Eurographics Association, pp. 1383–1392. URL: <http://dl.acm.org/citation.cfm?id=1735603.1735621>. 2
- [SPK*03] SUN Y., PAIK J., KOSCHAN A., PAGE D. L., ABIDI M. A.: Point fingerprint: a new 3-d object representation scheme. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 33, 4 (Aug 2003), 712–717. doi:10.1109/TSMCB.2003.814295. 2
- [SRGB14] SOLOMON J., RUSTAMOV R., GUIBAS L., BUTSCHER A.: Earth mover’s distances on discrete surfaces. *ACM Trans. Graph.* 33, 4 (July 2014), 67:1–67:12. doi:10.1145/2601097.2601175. 9
- [SSK*05] SURAZHISKY V., SURAZHISKY T., KIRSANOV D., GORTLER S. J., HOPPE H.: Fast exact and approximate geodesics on meshes. *ACM Trans. Graph.* 24, 3 (July 2005), 553–560. doi:10.1145/1073204.1073228. 4, 5
- [TBW*11] TEVS A., BERNER A., WAND M., IHRKE I., SEIDEL H.-P.: Intrinsic shape matching by planned landmark sampling. *Computer Graphics Forum* 30, 2 (2011), 543–552. URL: <http://dx.doi.org/10.1111/j.1467-8659.2011.01879.x>, doi:10.1111/j.1467-8659.2011.01879.x. 1, 2, 5, 10
- [YBS05] YOSHIZAWA S., BELYAEV A., SEIDEL H.-P.: Fast and robust detection of crest lines on meshes. In *Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling* (New York, NY, USA, 2005), SPM ’05, ACM, pp. 227–232. doi:10.1145/1060244.1060270. 8, 9