Intrinsic Shape Context Descriptors for Deformable Shapes

Iasonas Kokkinos *
Center for Visual Computing
Ecole Centrale Paris/INRIA Saclay
France

Michael M. Bronstein†
Faculty of Informatics
Università della Svizzera Italiana
Lugano, Switzerland

Roee Litman, Alex M. Bronstein‡
School of Electrical Engineering
Tel Aviv University
Israel

Abstract

In this work, we present intrinsic shape context (ISC) descriptors for 3D shapes. We generalize to surfaces the polar sampling of the image domain used in shape contexts: for this purpose, we chart the surface by shooting geodesic outwards from the point being analyzed; 'angle' is treated as tantamount to geodesic shooting direction, and radius as geodesic distance. To deal with orientation ambiguity, we exploit properties of the Fourier transform. Our charting method is intrinsic, i.e., invariant to isometric shape transformations. The resulting descriptor is a meta-descriptor that can be applied to any photometric or geometric property field defined on the shape, in particular, we can leverage recent developments in intrinsic shape analysis and construct ISC based on state-of-the-art dense shape descriptors such as heat kernel signatures. Our experiments demonstrate a notable improvement in shape matching on standard benchmarks.

1. Introduction

Over the past decade, feature-based methods have gained popularity in the computer vision and pattern recognition communities with the introduction of the scale invariant feature transform (SIFT) [28] and similar algorithms [2]. The ability of these methods to demonstrate sufficiently good performance in many settings, including object recognition and image retrieval and the public availability of the code made SIFT-like approaches a first choice and a *de facto* standard in a variety of image analysis tasks.

A significant milestone in the construction of this arsenal of tools was the *shape context* (SC) descriptor introduced in [3]. The original shape context approach was developed for planar shapes, and is based on the observation that a set of vectors connecting a point on the shape to the rest of the shape points constitutes a rich and discriminative description of the shape. In order to make this description manageable, the authors proposed to compute its distribution over a polar or log-polar system of coordinates, binning together vectors with similar length and orientation. In this way, each point on a shape is represented by a low-dimensional descriptor that aggregates information from distant locations. The original technique was subsequently extended to images [5, 31] where a point is described by the spatial distribution of some vector field containing appearance information (e.g., color). This has been observed [31] to boost matching and retrieval performance when compared to using descriptors based on solely local values of the field.

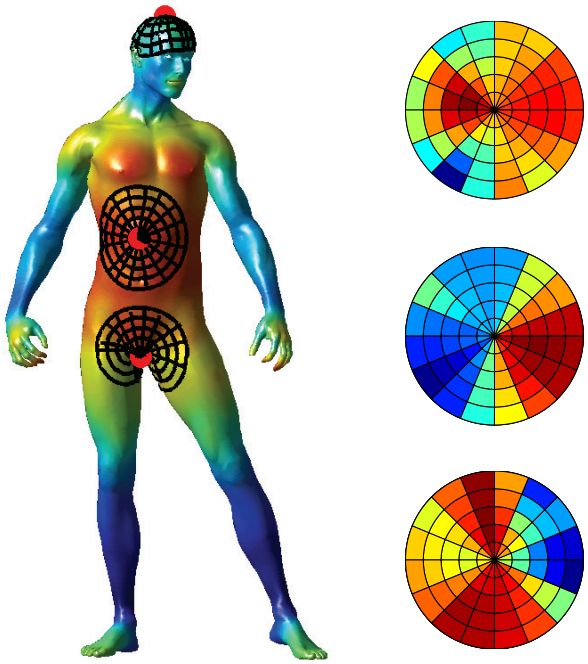


Figure 1. Our Intrinsic Shape Context (ISC, right) descriptor is constructed as a histogram of some field, defined on an intrinsic local polar coordinate system (left, shown at three different points).

Feature-based approaches have also been considered for

*Supported by ANR-10-JCJC -0205, † HP2C, ‡ ISF and GIF grants.

3D shapes and curved surfaces. Early research focused mainly on invariance under global Euclidean transformations (rigid motion). Classical works in this category include the integral volume descriptors [29, 17], spin images [21] and multiscale local features [32] just to mention a few out of many.

In the past decade, significant effort has been invested in extending the invariance properties to non-rigid deformations. Some of the classical rigid descriptors were extended to the non-rigid case by replacing the Euclidean metric with its geodesic counterpart [20, 14]. Also, the use of conformal factors has been proposed [27]. Being intrinsic properties of a surface, both are independent of the way the surface is embedded into the ambient Euclidean space and depend only on its metric structure. This makes such descriptors invariant to inelastic bending transformations. However, geodesic distances suffer from strong sensitivity to topological noise, while conformal factors, being a local quantity, are influenced by geometric noise. Both types of noise, virtually inevitable in real applications, limit the usefulness of such descriptors.

Recently, a family of intrinsic geometric properties broadly known as *diffusion geometry* has become growingly popular. The studies of diffusion geometry are based on the theoretical works by Berard *et al.* [4] and later by Coifman and Lafon [13] who suggested to use the eigenvalues and eigenvectors of the Laplace-Beltrami operator associated with the shape to construct invariant metrics known as diffusion distances. These distances as well as other diffusion geometric constructs [33, 10] have been shown significantly more robust compared to their geodesic counterparts. Diffusion geometry offers an intuitive interpretation of many shape properties in terms of spatial frequencies and allows to use standard harmonic analysis tools [26].

In their influential paper Sun *et al.* [37] (and independently, Gebal *et al.* [16]) introduced the *heat kernel signature* (HKS), based on the fundamental solutions of the heat equation (heat kernels). A scale-invariant version of HKS (SIHKS) was developed in [11], and a photometric HKS in [25]. In [1], another physically-inspired descriptor, the *wave kernel signature* (WKS) was proposed as a solution to the excessive sensitivity of the HKS to low-frequency information. As of today, these descriptors achieve state-of-the-art performance in many deformable shape analysis tasks such as shape retrieval [7] and 3D medical data analysis [12].

While being very successful in image analysis, feature based descriptors have so far achieved a more modest success in the analysis of surfaces. One of the difficulties in generalizing classical 2D techniques to non-Euclidean surfaces stems from the fact that unlike images, surfaces lack a global system of coordinates and can be associated with only a local vector structure (tangent plane). Such a system of coordinates holds only locally; trying to use it globally like done in spin images [21] makes the descriptor sensitive to shape deformations. Another difficulty is the fact that geometric information is usually much poorer in features, making it difficult to extract many repeatable and discriminative features as in images. This calls for methods that can operate everywhere on a surface, without relying on some special treatment, e.g. to determine local orientation.

Recently, there have been attempts to adapt popular image feature detectors and descriptors such as the Harris corner detector [36], difference of Gaussians and histogram of gradients [38] to 3D shapes. In particular, the latter work develops the gradient of a texture function plotted on a triangular mesh to create SIFT-like 3D descriptors. While dependent on triangulation and thus not purely intrinsic, this line of works in one of the inspirations of our paper.

One of the reasons for the success of the shape context descriptor is its "glocal" nature. We believe that contextual information can greatly improve the performance of existing surface descriptors. In this paper, we propose a generalization of the shape context approach to curved surfaces. As our field function we can use texture (appearance) information, or an intrinsic geometric signature such as HKS or SIHKS. Our goal, as in 2D SC, is to aggregate this information around the neighborhood of a point. The main difference is that we do it in an intrinsic manner that respects the surface geometry, and therefore call our descriptor *intrinsic shape context* (ISC).

3D Shape contexts were explored earlier in [24, 15]; but these works are not intrinsic, i.e. surface deformations affect the descriptors. An earlier work that explored the exploitation of intrinsic geometry was [35], but there the authors trivially deal with radial variation, by averaging over orientations; this is strictly subsumed by our approach, which allows us to retain and exploit the information contained in the radial variation around a point.

The contribution of this paper is three-fold. First, we develop a proper generalization of shape contexts to surfaces and show several ways to construct them numerically. Second, we show how to overcome the orientational ambiguity that unavoidably arises in the construction of local coordinate systems. Finally, we experimentally demonstrate that the introduction of spatial context significantly improves the discriminative power of the descriptor in matching and retrieval.

2. Shape Context Descriptors

The SC descriptor, originally developed for 2D shape analysis [34] and then extended to 2D image analysis in [5, 31], captures the values of a field $I(x)$, $x \in \mathbb{R}^2$ around a point x_i in a low-dimensional descriptor $\mathcal{S}(x_i) \in \mathbb{R}^{N_\rho \times N_\theta}$. SC constructs a polar grid with N_θ angular and N_ρ radial bins around x_i and forms the descriptor by averaging I over

each bin:

$$\mathcal{S}_{\rho,\theta}(x_i) \quad = \quad \frac{\int_{\mathbb{R}^2} \pi_{\rho,\theta}(x) I(x) dx}{\int_{\mathbb{R}^2} \pi_{\rho,\theta}(x) dx}, \tag{1}$$

where $\rho \in \{\rho_1, \ldots, \rho_{N_\rho}\}$ and $\theta \in \{\theta_1, \ldots, \theta_{N_\theta}\}$ are centers of radial and angular bins, respectively, and $\pi_{\rho,\theta} = \pi_\rho \pi_\theta$ denotes the membership function of angular bin θ and radial bin ρ,

$$\pi_\rho(x) \quad = \quad \begin{cases} 1 & \|x - x_i\|_2 \in R_\rho, \\ 0 & \text{else}, \end{cases} \tag{2}$$

$$\pi_\theta(x) \quad = \quad \begin{cases} 1 & \angle(x - x_i) \in R_\theta \\ 0 & \text{else}. \end{cases} \tag{3}$$

Here, $\|x - x_i\|_2$ is the radial and $\angle(x - x_i)$ the angular component of the polar coordinates system around x_i, and R_ρ, R_θ denote the supports of the radial and angular bins centered around ρ, θ, respectively. Thus, $\mathcal{S}_{\rho,\theta}$ is the average of $I(x)$ over the bin ρ, θ.

The hard binning in 2–3 can be replaced by soft binning using a different membership function that assigns to each point the probability to fall in a radial or angular bin.

3. Intrinsic Shape Context Descriptors

The main contribution of our work is the generalization of the shape context (SC) descriptors to surfaces. For this purpose, we adapt the polar (or log-polar) image sampling scheme used by SC to work with two-dimensional manifolds (surfaces). Our treatment is intrinsic and thereby invariant to isometric surface deformations and embedding of the surface in 3D space.

The construction of the polar bins of the SC descriptor in Eq.s 2–3 is valid only for images and cannot be applied straightforwardly to surfaces, for two reasons. First, such a construction assumes planar (zero-curvature) geometry, while 3D objects are generally curved. Second, images have a global unambiguous system of coordinates (allowing to define the angular coordinate $\angle(x - x_i)$) which surfaces usually do not have. Consequently, generalizing SC descriptors to surfaces involves three subproblems: charting the surface around a point, gathering statistics within the bins of the chart, and eliminating orientation ambiguity. We lay out our approach to these problems below.

3.1. Surface Charting

In order to build a shape context descriptor around a surface point, we need to establish a local polar grid. The main challenge here comes from the fact that the grid has to be intrinsic, since (as opposed to the case of images) surfaces have generally non-trivial curvature.

Let us be given a connected surface represented as a triangular mesh (X, E, T) with n vertices $X = \{x_1, \ldots, x_n\}$, N_E edges $E = \{(x_{i_1}, x_{j_1}), \ldots, (x_{i_{N_E}}, x_{j_{N_E}})\}$ s.t. $(x_i, x_j) \in E$ iff $(x_j, x_i) \in E$, and N_T triangular faces $T = \{(x_{i_1}, x_{j_1}, x_{k_1}), \ldots, (x_{i_{N_T}}, x_{j_{N_T}}, x_{k_{N_T}})\}$ s.t. (x_i, x_j), (x_j, x_k), $(x_i, x_k) \in E$. The mesh can be considered as a piece-wise planar approximation of the underlying smooth surface.

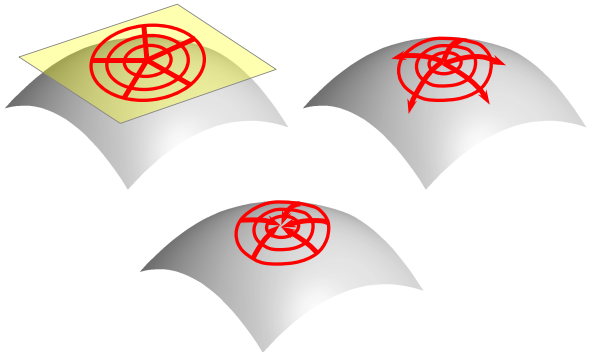


Figure 2. Three different possibilities of creating the intrinsic polar grid: using local mapping into the plane e.g. with MDS (top left), shooting geodesic directions from vertex (top right) and shooting directions from a geodesic circle (bottom). The approach adopted in this paper is the second one (top right).

Furthermore, we assume that the mesh is *manifold* (possibly with boundary), i.e., each edge is shared by at most two triangles. Edges belonging to one triangle are boundary edges, and corresponding vertices are *boundary points*. The surface around boundary points is topologically equivalent to a half-plane; around interior points, the surface is locally equivalent to a plane.

Finally, we denote by $d_X : X \times X \to \mathbb{R}_+$ the geodesic distance function, measuring the length of the shortest path on the mesh (not necessarily along the edges) between any pair of vertices.

Our goal is to create an intrinsic polar (or log-polar) grid around a vertex x_i. The smallest *1-ring* neighborhood of x_i is the set of all directly connected vertices, i.e. $N_1(x_i) = \{x_j : (x_i, x_j) \in E\}$. The creation of the grid can be done in three different ways.

Local MDS. One way is to find a planar representation $\xi : B_r(x_i) \to \mathbb{R}^2$ of a patch $B_r(x_i) = \{x_j \in X : d_X(x_i, x_j) \leq r\}$ around x_i, construct the grid on $\xi(B_r(x_i)) \subset \mathbb{R}^2$, and then map the grid back to the surface (Figure 2, top left). Such mappings can be established using *multidimensional scaling* (MDS) [6] or *conformal mapping* [19, 18, 39]. The former tries to create an isometric mapping from the patch to the plane by finding a planar configuration of points whose pairwise Euclidean distance are as close as possible to the pairwise geodesic distances between the surface points. However, neither MDS nor conformal map allows to precisely map the non-Euclidean geometry

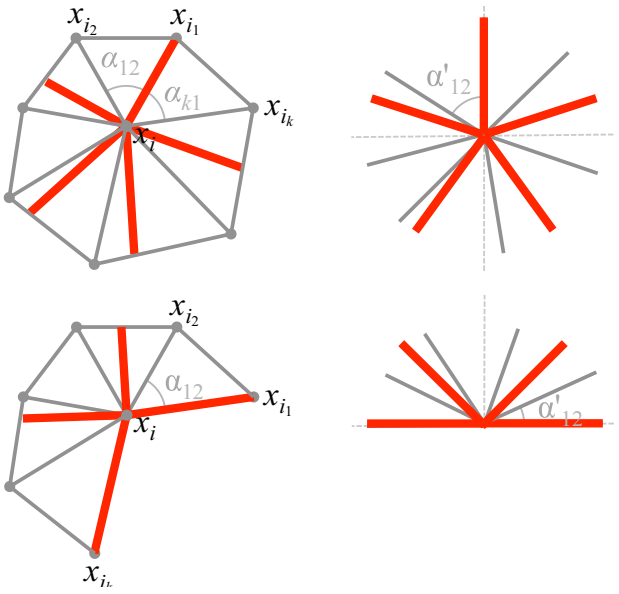


Figure 3. Creation of intrinsic angular chart at interior (top row) and boundary (middle row) point x_i. Because of surface curvature at x_i, the sum of angles $\Sigma\alpha = \alpha_{12} + \alpha_{23} + \ldots + \alpha_{k-1,k} + \alpha_{k1}$ in 1-ring of x_i is not necessarily equal to 2π. The partition is done by mapping the 1-ring into the plane (half plane in case of a boundary point) with angles $\alpha'_{ml} = \alpha_{ml}/\Sigma\alpha$, dividing the plane into equal angular segments, and mapping them back onto the 1-ring.

of the surface into the plane, thus resulting in angular or radial distortions.

Outward ray shooting. Another approach, adopted in this paper, is to construct the grid intrinsically (Figure 2, top right), in a procedure bearing resemblance to the fast marching method (FMM) used for geodesic distance computation [22]. In this approach, we chart the surface locally by shooting geodesics outwards from the vertex x_i. These directions provide us with the counterpart of rays in a log-polar mapping, i.e., they are surface loci with constant intrinsic angular coordinate. Along the same lines, we recover the counterpart of circles in a log-polar mapping by computing geodesic distances between the point x_i and all surface points x_j, and then recovering equidistant points (r-level sets of the function $d_X(x_i, x) = r$).

The initial directions are established by partitioning the 1-ring of an interior vertex x_i into segments of equal angle (Figure 3, top). This partition is performed as follows: we map the 1-ring triangles (Figure 3, top left) into the plane by partitioning the plane into angular segments (Figure 3, top right) whose angle rations are equal to the ratios of the angles in 1-ring triangles. Since the system of coordinates is local, there is a rotation ambiguity: one of the edges is chosen arbitrarily to align with the plane horizontal or vertical axis. The plane is divided into equiangular segments. The boundaries of these segments (denoted by red lines in Figure 3) are mapped back to the 1-ring triangles. In the limit this is equivalent to subdividing the tangent plane: flattening applies a common scaling to all surface triangles when mapping to the plane, so angle ratios on the plane will be preserved on the surface. For boundary vertices, the procedure is very similar, with the exception that the mapping is performed into a half-plane (Figure 3, bottom).

The propagation of the directions outwards from the 1-ring is done using the standard unfolding procedure [22, 8], which is known to be numerically consistent (refining the mesh will converge to the geodesic, for the same reason that fast marching converges). The main idea of unfolding is to create a poly-linear path by propagating a direction across adjacent triangles, as depicted in Figure 4.

Inward ray shooting. Finally, it is possible to create the angular partition by splitting a level curve of the geodesic distance function $d_X(x_i, x) = r$ into segments of equal length, and propagate directions from these points along the gradient $\nabla_X d_X(x_i, x)$ towards x_i (Figure 2, bottom).

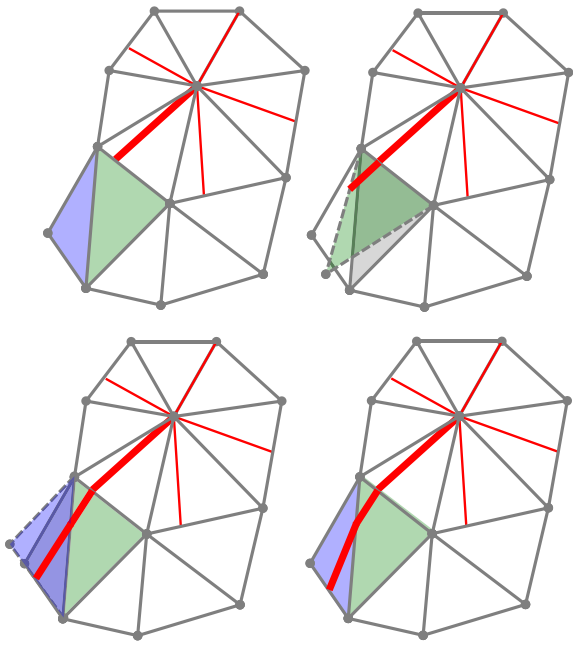


Figure 4. Propagation of direction outwards from 1-ring using the unfolding procedure. The green triangle adjacent to the 1-ring triangle is brought into the same plane as the one-ring triangle (shown in dashed, top right). The direction (thick red line) is continued until it hits the triangle edge. Then, the next adjacent triangle (blue) is brought into the same plane (bottom left). The procedure continues until the length of the resulting polyline reaches some threshold.

3.2. Surface-based Statistics

Constructing a shape context descriptor requires averaging the values of a field over radial and angular bins. We do

it similarly to the 2D shape context, replacing all the definitions in Eq. 1 by their intrinsic equivalents. The intrinsic shape context descriptor $\mathcal{S}_{\rho,\theta}(x_i)$ at point x_i is given by

$$\mathcal{S}_{\rho,\theta}(x_i) = \frac{\int_X \pi_{\rho,\theta}(x) I(x) d\mu_X(x)}{\int_X \pi_{\rho,\theta}(x) d\mu_X(x)}, \quad (4)$$

where

$$d\mu_X(x) = \frac{1}{3} \sum_{\substack{x_i,x_j \in N_1(x) \\ (x_i,x_j)\in E}} \text{area}(x, x_i, x_j)$$

is the local area element equal to one third the area of the 1-ring neighborhood of x, and $\pi_{\rho,\theta} = \pi_\rho \pi_\theta$, as previously, denotes the membership function of intrinsic angular bin θ and intrinsic radial bin ρ.

We use soft membership functions, constructed as follows. Considering that we use N_ρ radial bins, we perform a softmax assignment using the geodesic distance from x_i,

$$\pi_\rho(x) = \frac{\mu_\rho(d_X(x_i,x))}{\sum_\rho \mu_\rho(d_X(x_i,x))}, \quad (5)$$

$$\mu_\rho(r) = \begin{cases} \exp(\frac{\beta}{\rho^2}(r-\rho)^2) & \rho \le N_\rho - 1 \\ \frac{1}{1+\exp(-\gamma(x-\rho))} & \rho = N_\rho. \end{cases} \quad (6)$$

The membership function $\mu_\rho(r)$ indicates the extent to which radius r belongs to geodesic ring centered around ρ. If we use N_ρ rings, the first $N_\rho - 1$ functions are Gaussians, centered around ρ, while the last one is a sigmoidal, capturing the 'tail' of large radii. We set the parameters $\beta = 1/4, \gamma = 1$ so as to give a smooth transition among radii.

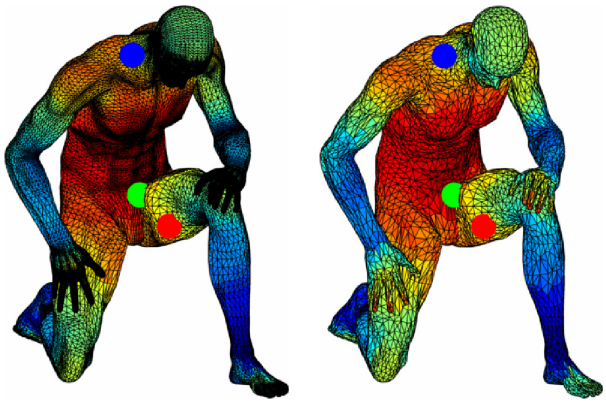
For the angular membership function, we denote by

$$d_{X,\theta}(x) = \frac{d_X(v(x_i,\theta),x)}{d_X(x_i,x)}$$

the normalized point-to-set geodesic distance between a point x and a polyline $v(x_i,\theta)$ shot from x_i in the direction θ (found as described in Section 3.1). The normalization guarantees that the softness of the assignment is not affected by the distance from x_i. The angular membership function is given by

$$\pi_\theta(x) = \frac{\exp(-\alpha d_{X,\theta}(x))}{\sum_\theta \exp(-\alpha d_{X,\theta}(x))}. \quad (7)$$

This softmax operation uses the distance of x to all candidate directions to estimate the 'probability' that x_i belongs to the angular bin θ. Due to normalization, we have $\pi_\theta(x) \in [0,1]$ and $\sum_\theta \pi_\theta(x) = 1$. The parameter α determines the hardness of the assignment and is set $\alpha = 10$.



52565 Mesh points 26299 Mesh Points

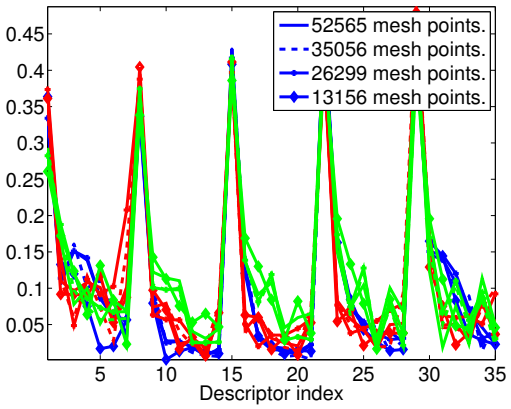


Figure 5. Effect of surface decimation on our descriptor: the color encodes the point, while the linestyle encodes the surface resolution at which the descriptor was computed. Our descriptor is largely invariant to decimation and remains discriminative.

3.3. Eliminating Orientation Ambiguity

Since our descriptor is intrinsic, it cannot rely on a global system of coordinates and constructs a local polar coordinate system around each vertex. This introduces rotation ambiguity: the ray shooting procedure outlined in Section 3.1 picks the starting orientation randomly, resulting in our ISC to be defined up to some unknown phase, $\mathcal{S}_{\rho,\theta+c \bmod 2\pi}$. It is possible in theory to determine some dominant reference orientation, but on surfaces -unlike images- we do not have intensity gradients, or similar cues to indicate orientations (for instance for a point on a sphere all orientations are equally good).

Here, we opt for using the *Fourier transform modulus* (FTM) technique to achieve rotation invariance. This technique has been used extensively in image registration [40], and was introduced for scale- and rotation- invariant descriptor construction in [23], and was also exploited to construct scale-invariant heat kernel signatures (SIHKS) for surfaces [11]. Here we exploit this technique to eliminate

the effects of orientation from our ISC descriptor.

We observe that transforming the descriptor in the Fourier domain w.r.t. the angular coordinate θ, we have

$$\mathcal{F}_\theta\{\mathcal{S}_{\rho,\theta}\}(\omega) = \sum_\theta \mathcal{S}_{\rho,\theta} \exp(-\iota\theta\omega) \quad (8)$$

$$\mathcal{F}_\theta\{\mathcal{S}_{\rho,\theta+c}\}(\omega) = \mathcal{F}_\theta\{\mathcal{S}_{\rho,\theta}\}(\omega) \exp(-\iota c\omega) \quad (9)$$

Taking the absolute value we have $|\mathcal{F}_\theta\{\mathcal{S}_{\rho,\theta+c}\}(\omega)| = |\mathcal{F}_\theta\{\mathcal{S}_{\rho,\theta}\}(\omega)|$, eliminating the effects of orientation ambiguity without relying on orientation selection.

The robustness and invariance of our descriptor is demonstrated in Fig. 5: we consider a range of surface decimation factors, and plot descriptors at 3 points which are common to all surfaces. We observe that the descriptor is largely invariant to the surface resolution -and also the associated changes in the 1-ring neighborhood- while at the same time discriminating among the three points. We now proceed to a thorough evaluation on standard benchmarks.

4. Experimental results

The performance of our descriptor was evaluated on the SHREC'10 robust correspondence benchmark [9]. The benchmark contains three distinct shape classes (human, dog, and horse), each shape undergoing ten different transformations (isometry, topology, sampling, global scaling, local scaling, holes, micro holes, Gaussian noise, and shot noise) with five strengths per transformation (from mild to very strong). Shapes are represented as triangular meshes with about 5×10^4 vertices (except for the sampling transformations, where the meshes are progressively decimated down to about 2.5×10^3 vertices). The benchmark also contained vertex-wise groundtruth correspondences between the transformed shapes and the reference (null) shapes, including intrinsic bilateral symmetries.

We created ISC based on HKS [37] and SIHKS [11] dense descriptors. These feature descriptors showed the best performance in SHREC benchmarks [9]. We used the cotangent weight scheme [30] to compute the first 300 eigenvalues and eigenvectors of the Laplace-Beltrami operator on each shape. Neumann boundary conditions were used. The HKS descriptors were computed at six time scales $t = 512, 675, 891, 1176, 1552$, and 2048. For the computation of the SIHKS, 400 logarithmically sampled scales from 10^{-5} to 5×10^5 were used. For the construction of the ISC descriptors, radius was set to $R = 20$; 5 linearly spaced radial bins and 16 angular bins were used. It takes 1 sec/vertex to compute an ISC descriptor using unoptimized Matlab code. Our construction involves geodesic distances from one vertex to a small neighborhood bounded by R, allowing us to terminate fast marching when reaching R.

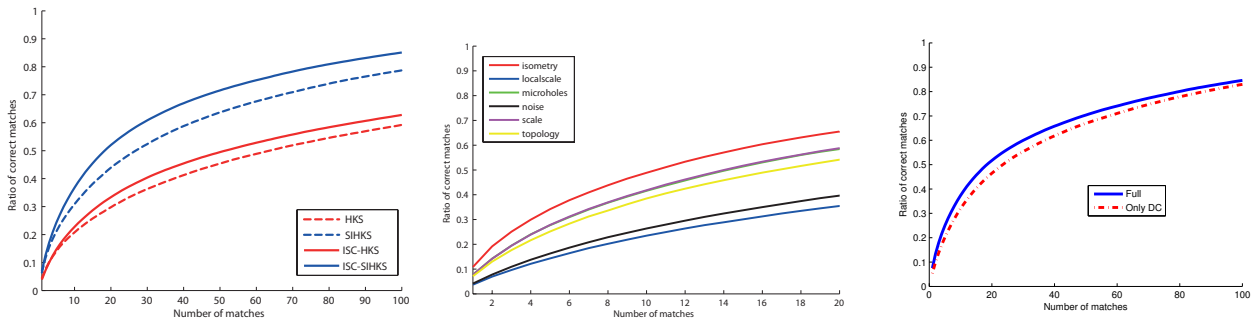
Figure 6 visualizes the distance maps computed in the descriptor space from a reference point on the human shape



Figure 6. Normalized Euclidean distance between the descriptor at a reference point on the right hand (white dot) and descriptors computed at rest of the points of the same shape and a distinct shape (bottom row) using the SIHKS descriptor (left) and the ISC-SIHKS (right). Observe how the introduction of spatial context improves the discriminativity and localization of the descriptor.

to the rest of the points on that shape, as well as to the points of the dog shape. Two phenomena are clearly visible: First, using the same base descriptor (SIHKS) the ISC gives significantly better feature *localization*, in the sense that the distance grows fast as one moves away from the reference point. Second, the ISC exhibits better *discriminativity* by assigning higher relative distances to the points of the distinct dog shape, while the raw SIHKS confuses between the reference point and some points on the dog's paws.

Figures 7–8 show the cumulative matching scores (CMC) of the raw HKS and SIHKS descriptors and the ISC descriptors based on them. 1000 points were sampled from each shape using farthest point sampling. Points from the transformed shape were matched to the null shape and ordered using L_2 distance between the corresponding descriptors. Correct matches were defined as matches falling within 2% of the shape intrinsic diameter off the ground truth match (including the bilaterally symmetric one). This resulted in about 10^6 wrong matches and 7×10^3 correct ones per shape. The CMC curve shows the percentage of feature points that had correct match among the first k candidate matches. Note that the first 20 matches capture a substantial amount of correct matches, often higher than 50%.



a b c

Figure 7. Cumulative matching score curves (CMC) averaged over all types of shape transformations for (a) different descriptors, (b) the ISC-SIHKS descriptor and different types of shape transformations, (c) the full-blown ISC-SIHKS descriptor versus using only the DC components.

The use of spatial contexts consistently improves the descriptor performance under all transformations, on the average by over 10%. Moreover, the use of exclusively the DC components of the Fourier Transform -i.e. the result of averaging over orientations- harms performance when compared to using the full-blown descriptor.

5. Conclusion

We developed an intrinsic version of the popular shape context descriptor, making it suitable for the analysis of deformable 3D shapes. Our construction is based on local charting of the surface and creating of semi-local histogram of a vector field representing some geometric or photometric property of the shape. The intrinsic shape context is thus a "glocal" meta-descriptor that can be used in combination with any dense shape descriptor, e.g. heat kernel signature or wave kernel signature. The use of semi-local support increases the descriptor localization and discriminativity over HKS-type descriptors which only encode local information. Using the local system of coordinates, intrinsic versions of many other semi-local descriptors popular in image analysis, such as covariance descriptors, can be designed. We intend to explore these constructions in our future work.

References

[1] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature-a quantum mechanical approach to shape analyis. In *Proc. CVPR*, 2011. 2

[2] H. Bay, T. Tuytelaars, and L. Van Gool. SURF: Speeded up robust features. In *Proc. ECCV*, 2006. 1

[3] S. Belongie, J. Malik, and J. Puzicha. Shape context: A new descriptor for shape matching and object recognition. In *Proc. NIPS*, 2000. 1

[4] P. Bérard, G. Besson, and S. Gallot. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, 4(4):373–398, 1994. 2

[5] A. Berg and J. Malik. Geometric blur for template matching. In *CVPR*, 2001. 1, 2

[6] I. Borg and P. Groenen. *Modern multidimensional scaling - theory and applications*. Springer, 1997. 3

[7] A. M. Bronstein, M. M. Bronstein, L. J. Guibas, and M. Ovsjanikov. Shape google: geometric words and expressions for invariant shape retrieval. *TOG*, 30(1):1, 2011. 2

[8] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Scientific Computing*, 28(5):1812–1836, 2006. 4

[9] A. M. Bronstein et al. SHREC 2010: robust correspondence benchmark. In *Proc. 3DOR*, 2010. 6

[10] M. M. Bronstein and A. M. Bronstein. Shape recognition with spectral distances",. *PAMI*, 33(5):1065–1071. 2

[11] M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR*, 2010. 2, 5, 6

[12] U. Castellani, P. Mirtuono, V. Murino, M. Bellani, G. Rambaldelli, M. Tansella, and P. Brambilla. A new shape diffusion descriptor for brain classification. In *Proc. MICCAI*, 2011. 2

[13] R.R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. 2

[14] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *PAMI*, pages 1285–1311, 2003. 2

[15] A. Frome, D. Huber, R. Kolluri, T. Bülow, and J. Malik. Recognizing objects in range data using regional point descriptors. In *Proc. ECCV*, 2004. 2

[16] K. Gebal, J.A. Bærentzen, H. Aanæs, and R. Larsen. Shape analysis using the auto diffusion function. In *Computer Graphics Forum*, volume 28, pages 1405–1413, 2009. 2

[17] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global registration. In *Proc. SGP*, 2005. 2

[18] X. Gu and S.-T. Yau. Global Conformal Surface Parameterization. In *Proc. SGP*, 2003. 3

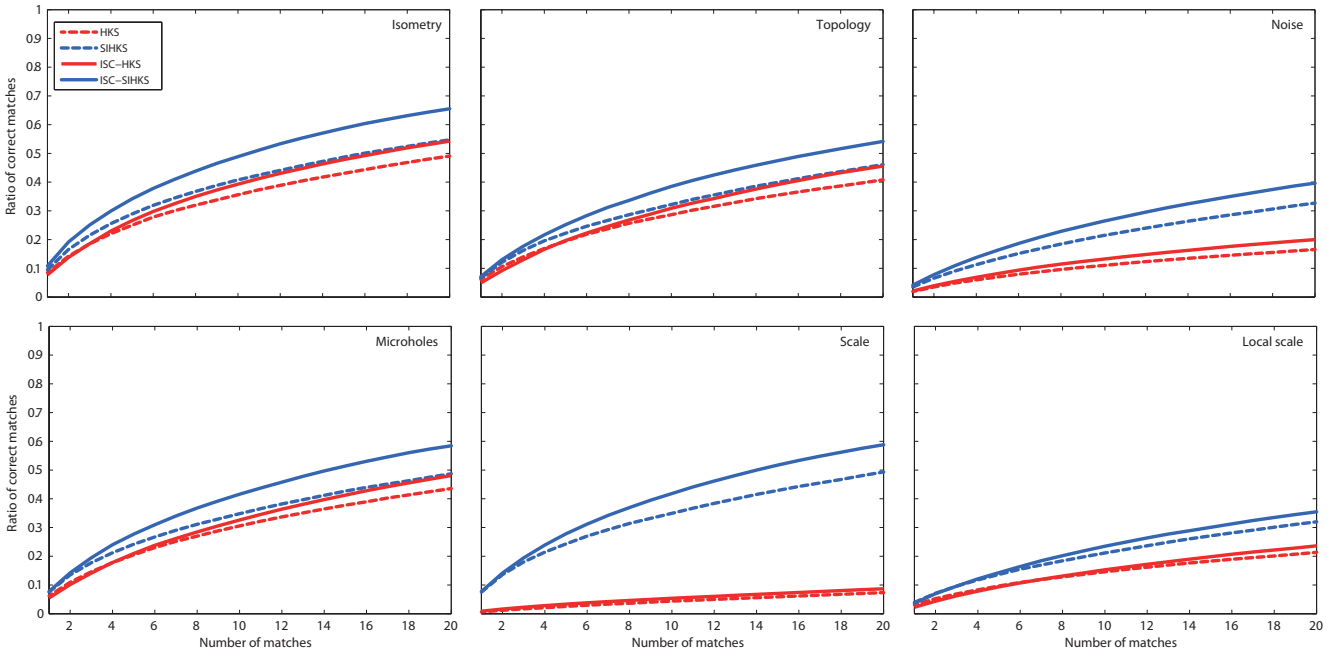


Figure 8. CMC curves of different shape descriptors broken down by different types of shape transformations.

[19] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, and M. Halle. Conformal surface parameterization for texture mapping. *PAMI*, 2000. 3

[20] A. B. Hamza and H. Krim. Geodesic object representation and recognition. In *Discrete Geometry for Computer Imagery*, pages 378–387, 2003. 2

[21] A. E. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3D scenes. *PAMI*, 21(5):433–449, 1999. 2

[22] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *PNAS*, 95(15):8431–8435, 1998. 4

[23] I. Kokkinos and A. Yuille. Scale Invariance without Scale Selection. In *CVPR*, 2008. 5

[24] M. Körtgen, G. J. Park, M. Novotni, and R. Klein. 3d shape matching with 3d shape contexts. In *The 7th Central European Seminar on Computer Graphics*, 2003. 2

[25] A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, and R. Kimmel. Photometric heat kernel signature. In *Proc. SSVM*, 2011. 2

[26] B. Lévy. Laplace-Beltrami eigenfunctions towards an algorithm that understands geometry. In *Proc. SMI*, 2006. 2

[27] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. In *TOG*, 2009. 2

[28] D. Lowe. Distinctive image features from scale-invariant keypoint. *IJCV*, 60(2):91–110, 2004. 1

[29] S. Manay, B.W. Hong, A.J. Yezzi, and S. Soatto. Integral invariant signatures. In *Proc. ECCV*, 2004. 2

[30] M. Meyer, M. Desbrun, P. Schroder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and Mathematics III*, pages 35–57, 2003. 6

[31] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *PAMI*, 2005. 1, 2

[32] M. Pauly, R. Keiser, and M. Gross. Multi-scale feature extraction on point-sampled surfaces. In *Computer Graphics Forum*, volume 22, pages 281–289, 2003. 2

[33] R.M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP*, 2007. 2

[34] S. Belongie and J. Malik and J. Puzicha. Shape matching and object recognition using shape contexts. *PAMI*, 2002. 2

[35] Y. Shi, P. Thompson, G. Zubicaray, S. Rose, Z. Tu, I. Dinov, and A. Toga. Direct mapping of hippocampal surfaces with intrinsic shape context. *Neuroimage*, 37(3):792–807, 2007. 2

[36] I. Sipiran and B. Bustos. A robust 3D interest points detector based on Harris operator. In *Proc. 3DOR*, 2010. 2

[37] J. Sun, M. Ovsjanikov, and L. J. Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. In *Computer Graphics Forum*, 2009. 2, 6

[38] A. Zaharescu, E. Boyer, K. Varanasi, and R Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. CVPR*, 2009. 2

[39] W. Zeng, D. Samaras, and X. Gu. Ricci flow for 3D shape analysis. *PAMI*, 32, 2010. 3

[40] S. Zokai and G. Wolberg. Image registration using log-polar mappings for recovery of large-scale similarity and projective transformations. *IEEE Trans. Im. Proc.*, 14, 2005. 5