

Scale-invariant heat kernel signatures for non-rigid shape recognition

Michael M. Bronstein

Department of Computer Science, Technion – Israel Institute of Technology
Haifa 32000, Israel

mbron@cs.technion.ac.il

Iasonas Kokkinos

Department of Applied Mathematics, École Centrale de Paris
Grande Voie des Vignes, 92295 Chatenay-Malabry, France

iasonas.kokkinos@ecp.fr

Abstract

One of the biggest challenges in non-rigid shape retrieval and comparison is the design of a shape descriptor that would maintain invariance under a wide class of transformations the shape can undergo. Recently, heat kernel signature was introduced as an intrinsic local shape descriptor based on diffusion scale-space analysis. In this paper, we develop a scale-invariant version of the heat kernel descriptor. Our construction is based on a logarithmically sampled scale-space in which shape scaling corresponds, up to a multiplicative constant, to a translation. This translation is undone using the magnitude of the Fourier transform. The proposed scale-invariant local descriptors can be used in the bag-of-features framework for shape retrieval in the presence of transformations such as isometric deformations, missing data, topological noise, and global and local scaling. We get significant performance improvement over state-of-the-art algorithms on recently established non-rigid shape retrieval benchmarks.

1. Introduction

Today, only a small fraction of Internet repositories of visual and geometric data is tagged and accessible through simple text search. Fast growth of these repositories makes content-based retrieval one of the next grand challenges in search and organization of such information. Particularly difficult is the problem of *shape retrieval*, as geometric shapes manifest a vast variability due to different scale, orientation, non-rigid deformations, missing data, and also appear in a variety of different formats and representations.

In principle, the common denominator of shape retrieval approaches is the creation of a *shape descriptor* or *signature*

which captures the unique properties of the shape that distinguish it from shapes belonging to other classes on the one hand, and is invariant to a certain class of transformations a shape can undergo on the other [40, 39]. In rigid shape analysis, different types of invariance were addressed. Rotation and translation invariance can be achieved using volume and area descriptors [45], spherical harmonics [16], geometric moments *et al.* [38], and distribution of pair-wise Euclidean distances [29].

Dealing with non-rigid shapes requires compensating for the degrees of freedom resulting from deformations. Elad and Kimmel [12] and follow-up works [26, 8] proposed modeling shapes as metric spaces with intrinsic (e.g. geodesic) distances, which are invariant to inelastic deformations. Ling and Jacobs [22] and Bronstein *et al.* [7] used this framework with a metric defined by internal distances in 2D shapes. Reuter *et al.* [33, 32] used the Laplacian spectra as intrinsic shape descriptors.

A particular type of intrinsic geometry is generated by heat diffusion processes on the shape. Coifman and Lafon [11] popularized the notion of *diffusion geometry*, which is closely related to scale-space methods in image processing [36]. Rustomov [34] was one of the first to use such distances in shape analysis, applying the method of Osada *et al.* [29] to *commute time distances* (this method is similar to the recent work of Mahmoudi and Sapiro [24] who used *diffusion distances* instead). In [5], shapes were analyzed as metric spaces equipped with diffusion metrics.

In image analysis, bottom-up approaches have become popular, notably due to the works of Zisserman *et al.* [35, 10] and Schmid *et al.* [27]. Using these approaches, an image is described as a collection of local features (“visual words”) from a given vocabulary, resulting in a representation referred to as a *bag of features*. In shape analysis, such approaches have been introduced more recently by

Ovsjanikov *et al.* [6] and Toldo *et al.* [42] (see [28, 20] for earlier similar ideas).

The bag of features paradigm relies heavily on the choice of the local feature descriptor that is used to create the visual words. In image analysis and 2D shape retrieval, typical features are blobs [25] and corners [27], and the default choice for a local descriptor is the scale-invariant feature transform (SIFT) [23] or one of its varieties [2, 41]. Scale-invariant local descriptors can be constructed in two ways. First way is to use scale-space analysis of the image to locally estimate the scale [21, 23]. Descriptors are then extracted from appropriately scaled image patches. Second way is to use a combination of logarithmic sampling with Fourier analysis to compensate for the scaling effects [17] (such an approach is also commonly used to compute a global image rotation and scaling in the context of registration [9, 46]).

In shape analysis, on the other hand, there is no commonly agreed upon feature descriptor similar to SIFT. In non-rigid shape retrieval applications, an ideal feature descriptor should be first of all intrinsic and thus deformation-invariant. Second, it should cope with missing parts, and also be insensitive to topological noise and connectivity changes. Third, it should work across different shape representations and formats (e.g. point clouds and meshes) and be insensitive to sampling. Finally, the descriptor should be scale-invariant. The last two properties are especially important when dealing with shapes coming from Internet repositories such as Google 3D Warehouse, where shapes appear in a variety of representations and with arbitrary scales.

Different approaches such as contour and edge features [30, 18], spin images [14], local patches [28, 42], conformal factor [4], differential operators [44], and local volume properties [13] were used as feature descriptors in shape retrieval literature. Unfortunately, none of them satisfy all of the above desired properties (for example, volumetric and patch-based methods are not intrinsic, and conformal factor is sensitive to topology).

Recently, a local feature descriptor based on multiscale heat kernels was proposed [37]. This descriptor satisfies all of the above properties except for scale invariance. Scale invariance poses an additional challenge, for a few reasons. Compared to images, shapes typically contain less features that would be roughly analogous to blobs or corners, and there is no clear generalization of such structures to 3D surfaces. Feature detection based on intrinsic scale-space analysis such as [37] would find a few reliable points (usually with high curvature), at which scale estimation can be done. In flat regions, no scale estimation is possible. For this reason, Ovsjanikov *et al.* [6] avoided feature detection and used a *dense* feature descriptor computed at every point of the shape in combination with statistical weighting to re-

duce the influence of trivial points.

Contribution. In this paper, we develop a scale-invariant version of the heat kernel signature by combining this descriptor with the recent approach of [17] to scale invariance in images. Our construction is based on a logarithmically sampled scale-space in which shape scaling corresponds, up to a multiplicative constant, to a translation. This translation is then undone using the magnitude of the Fourier transform. Since our descriptor does not rely on local scale estimation, it is computable at every point including flat regions, and can be thus used in the shape retrieval framework of [6], as well as for other applications such as dense correspondence between shapes.

2. Background

In the following discussion, we model shapes as Riemannian manifolds (possibly with boundary) and use the heat conduction properties as shape descriptors. Heat propagation on non-Euclidean domains is governed by the *heat diffusion equation*,

$$\left(\Delta_X + \frac{\partial}{\partial t}\right)u = 0, \quad (1)$$

where, Δ_X denotes the positive semi-definite *Laplace-Beltrami operator*, a Riemannian equivalent of the the Laplacian. The solution $u(x, t)$ of the heat equation with the initial conditions $u(x, 0) = u_0(x)$ (and respective boundary conditions if X has a boundary) describes the amount of heat on the surface at point x in time t . The solution of (1) with point heat distribution $u_0(x) = \delta(x - z)$ as initial conditions is called the *heat kernel* and denoted by $K_{X,t}(x, z)$.

On compact manifolds, the heat kernel can be presented as [15]

$$K_{X,t}(x, z) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x) \phi_i(z). \quad (2)$$

where $\lambda_0, \lambda_1, \dots \geq 0$ are eigenvalues and ϕ_0, ϕ_1, \dots are the corresponding eigenfunctions of the Laplace-Beltrami operator, satisfying $\Delta_X \phi_i = \lambda_i \phi_i$.

Heat kernel signatures. Sun *et al.* [37] proposed using the *heat kernel signature* (HKS)

$$h(x, t) = K_{X,t}(x, x) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i^2(x) \quad (3)$$

as local shape descriptors. The HKS is intrinsic and thus isometry-invariant (two isometric shapes have equal HKS), multi-scale and thus capture both local features and global shape structure, and also informative: under mild conditions, if two shapes have equal heat kernel signatures, they are isometric [37].

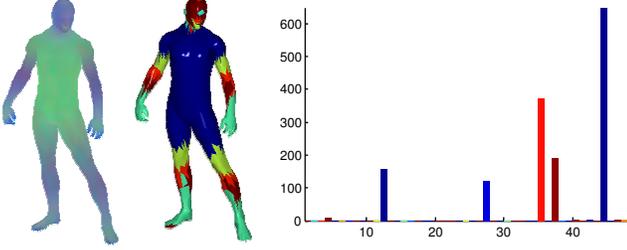


Figure 1. Construction of a bag of features shape descriptor. Left: dense HKS local descriptor (shown three components as RGB colors); middle: local descriptor quantized in a geometric vocabulary of size 48 (each color represents a geometric word); right: bag of features counting the frequency of appearance of each geometric word.

Ovsjanikov *et al.* [6] used the HKS to construct *global* shape descriptors following the bag of features paradigm used in image retrieval applications [35, 10]. First, the HKS descriptor is computed at every point of the shape (Figure 1, left). Next, using vector quantization, for each point on the shape, the HKS is replaced by the index of the most similar entry in a *geometric vocabulary* consisting of representative heat kernel signatures or “geometric words” (Figure 1, middle). The vocabulary is constructed offline by performing clustering in the HKS space. Finally, the distribution of geometric words on the shape is computed, resulting in a bag of features representation (Figure 1, right).

Sensitivity to scale. A notable disadvantage of the heat kernel signatures is their sensitivity to scale. Given a shape X and its scaled version $X' = \beta X$, the new eigenvalues and eigenfunctions will satisfy $\lambda' = \beta^2 \lambda$ and $\phi' = \beta \phi$. We therefore have the following equation:

$$h'(x, t) = \sum_{i=0}^{\infty} e^{-\lambda_i \beta^2 t} \phi_i^2(x) \beta^2 = \beta^2 h(x, \beta^2 t), \quad (4)$$

relating the signature h' at time t for X' with the signature h at time $\beta^2 t$ for X .

In some cases, the scaling effect can be undone using some *global pre-normalization* of the shape. Possible ways are normalizing the bounding box of the shape or its covariance (geometric moments), normalize the intrinsic diameter of the shape, i.e. the longest geodesic distance, or normalize the Laplace-Beltrami eigenvalues. The first approach will work only in rigid shapes, as non-rigid deformations change the bounding box. The second and the third approaches are insensitive to deformations, but would fail if the shape has missing parts. In the following, we describe an approach for *local* normalization of the heat kernel signature, which does not suffer from this problem.

3. Scale-invariant heat kernel signatures

In order to achieve scale invariance, we need to remove the dependence of h from the scale factor β . This is possible through the following series of transformations applied to h . First, at each shape point x we sample the heat signature logarithmically in time ($t = \alpha^\tau$) and form the discrete function

$$h_\tau = h(x, \alpha^\tau). \quad (5)$$

Based on Eq. 4, scaling the shape by β will result in a time-shift by $s = 2 \log_\alpha \beta$ and amplitude-scaling by β^2 (Figure 2, left):

$$h'_\tau = \beta^2 h_{\tau+s}. \quad (6)$$

Second, we remove the multiplicative constant β^2 by taking the logarithm of h , and then the discrete derivative w.r.t. to τ (Figure 2, middle). The first step turns the multiplicative factor into an additive constant, $2 \log \beta$, which then vanishes in differentiation:

$$\dot{h}'_\tau = \dot{h}_{\tau+s}, \quad (7)$$

(here, $\dot{h}_\tau = \log h_{\tau+1} - \log h_\tau$).

Finally, taking the discrete-time Fourier transform of \dot{h}_τ turns this shift in time into a complex phase;

$$H'(\omega) = H(\omega) e^{2\pi\omega s}, \quad (8)$$

where H and H' denote the Fourier transform of \dot{h} and \dot{h}' , respectively, and $\omega \in [0, 2\pi]$. The phase is in turn eliminated by taking the Fourier transform modulus (FTM):

$$|H'(\omega)| = |H(\omega)|. \quad (9)$$

We thus have constructed the scale-invariant quantity $|H(\omega)|$ (denoted as SI-HKS and shown in Figure 2, right) from the HKS at each point x , without performing scale selection. This allows us to compute descriptors at any point of our shape, where scale selection based on maxima detection could be impossible. Moreover, most of the signal information is contained in the low-frequency components of the FT, so we can build a compact descriptor by sampling $|H(\omega)|$ at a small number of low frequencies.

One caveat of our approach could be that scaling the shape and then resampling the function \dot{h}_τ makes the samples at the boundaries of the range of τ change. This can have dramatic effects if the signal information is concentrated at the boundaries of the scale-space. Fortunately, the HKS is typically smooth at low- and high- scales and therefore its derivative is equal to zero for a broad range of τs at the beginning and end of \dot{h} .

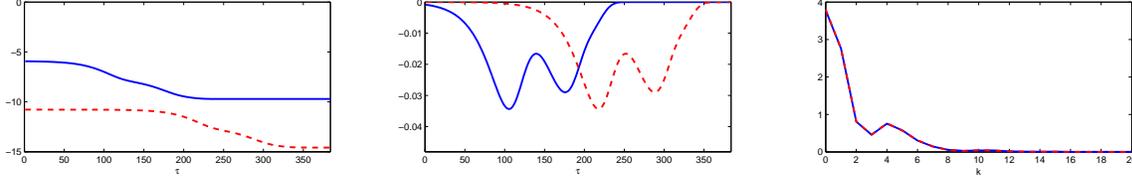


Figure 2. Construction of the scale-invariant heat kernel signature. Left: heat kernel signatures h (red) and h' (blue) computed at a corresponding point on a shape and its version scaled by the factor of 11, plotted on a logarithmic scale. h and h' differ by scale and shift in τ . Middle: h_τ and h'_τ , where the multiplicative constant is undone and the change in scale corresponds to a shift in τ only. Right: first 10 frequencies of $|H(\omega)|$ and $|H'(\omega)|$ used as scale-invariant HKS; the two descriptors computed at the two different scales are virtually identical.

4. Numerical computation

Numerical computation of the HKS and the SI-HKS is done using formula (2), in which a finite number of terms is taken and the continuous eigenfunctions and eigenvalues of the Laplace-Beltrami operator are replaced by the discrete counterparts. The discretization of the Laplace-Beltrami operator depends on the representation on the shape. For shapes represented as point clouds, the Laplace-Beltrami operator can be approximated using [3]. For triangular meshes, one of the most common discretizations is the *cotangent weight* scheme [31], defined for any function f on the mesh vertices as

$$(\Delta_{\hat{x}} f)_i = \frac{1}{a_i} \sum_j w_{ij} (f_i - f_j), \quad (10)$$

where $w_{ij} = \cot \alpha_{ij} + \cot \beta_{ij}$ for j in the 1-ring neighborhood of vertex i and zero otherwise (α_{ij} and β_{ij} are the two angles opposite to the edge between vertices i and j in the two triangles sharing the edge), and a_i are normalization coefficients proportional to the area of triangles sharing the vertex x_i . This discretization preserves many important properties of the continuous Laplace-Beltrami operator, such as positive semi-definiteness, symmetry, and locality, and in addition it is numerically consistent [43]. In matrix notation, Eq. (10) can be written as

$$\Delta_{\hat{x}} f = A^{-1} L f, \quad (11)$$

where $A = \text{diag}(a_i)$ and $L = \text{diag}(\sum_{l \neq i} w_{il}) - (w_{ij})$.

The eigenvalues and eigenfunctions of the Laplace-Beltrami operator discretized according to 11 are computed by solving the *generalized eigendecomposition* problem [19]

$$A \Phi = \Lambda L \Phi, \quad (12)$$

where Λ is the $(k+1) \times (k+1)$ diagonal matrix of the smallest eigenvalues $\lambda_0, \dots, \lambda_k$, and Φ is an $N \times (k+1)$ matrix of corresponding eigenfunctions ϕ_0, \dots, ϕ_k such that ϕ_{il} is the value of the l th eigenfunction at the point x_i . Another way of approximating Laplace-Beltrami eigenfunctions on

triangular meshes is using finite element methods (FEM) [32].

The discrete heat kernel signature is approximated by

$$h(x_l, \tau) \approx \sum_{l=0}^k e^{-\lambda_l \alpha^\tau} \phi_{il}^2 = \Psi e^{-T\Lambda}, \quad (13)$$

where $T = \text{diag}(\alpha^\tau)$ and $\Psi = (\phi_{il}^2)$. Since the heat kernel depends only on the eigenfunctions and eigenvalues of the Laplace-Beltrami operator, at least in theory, one can compare shapes in different representations (e.g., point clouds to meshes). This property of heat kernel signatures is especially appealing in Internet shape retrieval applications, where the variety of shape representations and formats is enormous.

5. Results

We used the ShapeGoogle database [6], consisting of 1061 shapes with simulated transformations. As of today, this is the largest non-rigid shape retrieval benchmark available. The database contained shapes from 469 different classes. For thirteen shape classes, the following transformations were simulated: 208 isometry, 208 global scale (varying approximately between 0.7 and 1.35), 128 local scale (local “swelling” of the shape), and 48 partiality+scale (missing parts in shapes with different global scaling). Examples of transformations are shown in Figure 4.

Heat kernel signatures (HKS) and the proposed scale-invariant heat kernel signatures (SI-HKS), respectively, were used as local shape descriptors. For the discrete computation of the heat kernels, we used the cotangent weight approximation of the Laplace-Beltrami operator and $k = 200$. For HKS, we used the parameters as in [6] (six scales 1024, 1351, 1783, 2353, 3104 and 4096), which were experimentally found to give optimal performance on the ShapeGoogle database. In order to construct the SI-HKS, we used a logarithmic scale-space with base $\alpha = 2$ and τ ranging from 1 to 25 with increments of $1/16$. After applying the logarithm, derivative, and Fourier transform, the first 6 discrete lowest frequencies were used as the local descriptor.

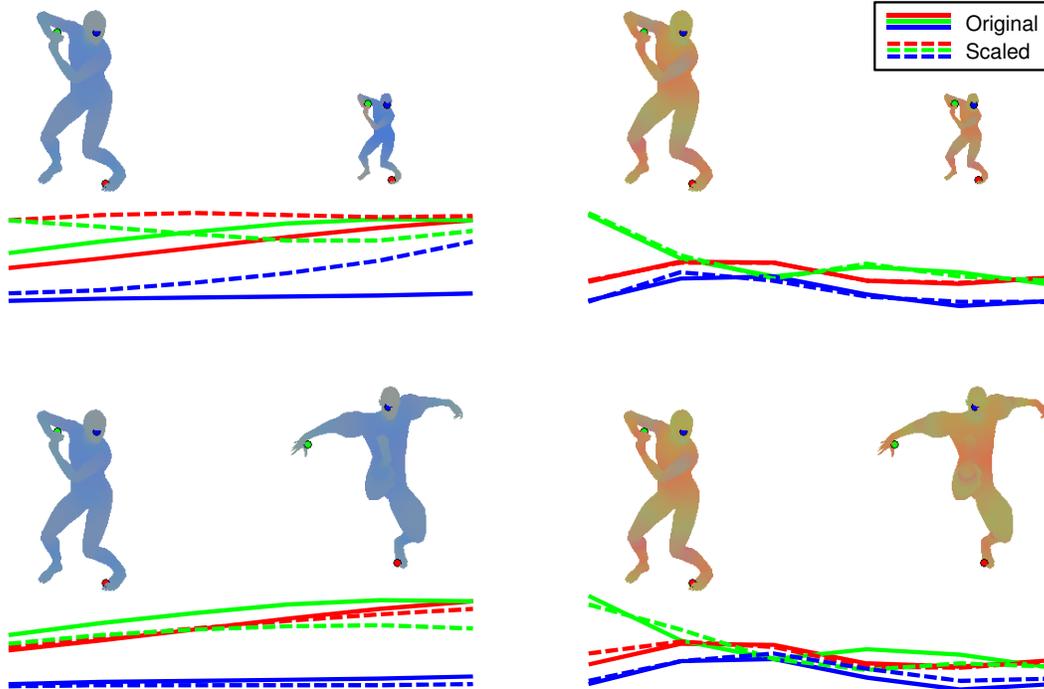


Figure 3. Comparison of HKS (left) and the proposed scale-invariant HKS (right). First and third rows: three components of HKS and SI-HKS, represented as RGB color and shown for shapes differing by global (first row) and local (third row) transformations. Second and fourth rows: HKS and SI-HKS at three point on the head (blue), hand (green), and foot (red) of the human shape.

Shape descriptors were constructed using bags of geometric words proposed in [6]. For HKS and SI-HKS, a geometric vocabulary of size 48 was built using clustering in the signature space (six-dimensional in both cases). The HKS and SI-HKS at each point of the shape were replaced by the closest geometric word from the vocabulary using soft vector quantization. We used the approximate nearest neighbor algorithm [1] as implemented in the ANN toolbox.¹ The distribution of geometric words (48-dimensional bag of features) was used as the shape descriptor. L_1 distance was used to compare the bags of features.

For comparison, we show the results of the ShapeDNA approach [33], describing shapes by the vector of the first eigenvalues of the Laplace-Beltrami operator. We used first 15 eigenvalues to construct the ShapeDNA descriptors (this parameter was empirically selected to achieve optimal performance on the ShapeGoogle database). Eigenvalues were computed using the same cotangent weight discretization. L_2 distance was used to compare the ShapeDNA descriptors.

Shape retrieval performance was quantified using the *precision-recall* (PR) curve (Figure 5), plotting the tradeoff between *precision* (ratio of the number of relevant shapes retrieved and the total number of shapes retrieved) and *re-*

call (ratio of the number of relevant shapes retrieved and the total number of existing relevant shapes that could be ideally retrieved). We used the *mean average precision* (mAP) as a single number to quantify the retrieval quality (average precision is computed as the area below the precision-recall curve for each query, and the mAP is the average of AP over all queries).

Table 1 shows the performance of shape retrieval using bags of features built of HKS and SI-HKS. Our approach shows a dramatic improvement in the presence of varying scale (99.5% mAP compared to 61.32% with HKS) and also better performance for local scaling transformations (92.60% mAP compared to 85.83%). HKS-based bags of features produce negligibly (by 0.01%) worse results than SI-HKS on the class of isometric deformations. ShapeDNA shows similar nearly perfect performance on the class of isometries, but performs very poorly on scale and local scale transformations (36.72% and 72.17% mAP, respectively)

Figure 6 shows examples of first five matches retrieved using HKS and SI-HKS. With HKS, a scaled down centaur is confused with a dog (row a) and a scaled up horse is confused with an elephant (row c); while SI-HKS produces correct matches (rows b and d). In the presence of local scaling, because of the local nature of the descriptor, it remains unchanged far from the deformed parts. The SI-HKS shows better robustness to such scaling compared to HKS (e.g., in

¹Code available from <http://www.cs.umd.edu/~mount/ANN>

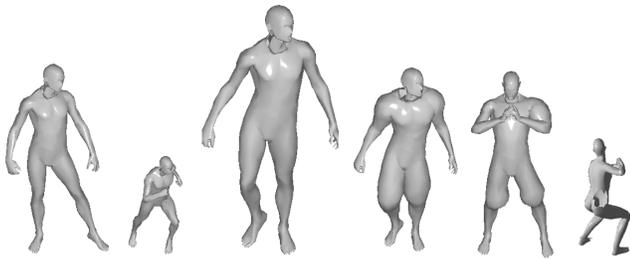


Figure 4. Example of transformations used in our shape retrieval experiment (left to right): null, scale down, scale up, two examples of local scale, partiality+scale.

Table 1. Shape retrieval performance (mAP in percents) using HKS and SI-HKS based bags of features and ShapeDNA [33]. Best result is shown in bold.

Transformation	Queries	HKS	SI-HKS	ShapeDNA
Isometry	208	99.96%	99.97%	99.52%
Scale	208	61.32%	99.95%	36.72%
Local scale	128	85.83%	92.60%	72.17%
Partiality+scale	48	54.67%	89.95%	27.42%
All	1061	85.30%	97.25%	74.47%

Figure 6 (g) local scaling transformations make HKS confuse between male and female shapes).

6. Conclusions

We presented an extension of the heat kernel signature allowing to deal with global and local scaling transformations. The use of Fourier transform magnitude to extract a scale-invariant quantity out of the heat kernel signature is advantageous over attempts to perform scale localization, which works only at prominent feature points. Our approach allows to create a dense scale-invariant feature descriptor defined at every point of the shape. Besides invariance to global scaling, the scale-invariant HKS shows better resilience to local scaling transformations. Such transformations can arise, for example, due to locally-elastic deformations that stretch or shrink the shape surface. In future work, we intend to explore the proposed method in the context of part-based approaches in which the local descriptor is confined to a part of a shape and not computed across parts, and a separate descriptor is computed for each part. This way, inelastic deformations could be addressed.

References

[1] S. Arya, D. Mount, N. Netanyahu, R. Silverman, and A. Wu. An optimal algorithm for approximate nearest neighbor

searching fixed dimensions. *Journal of the ACM*, 45(6):891–923, 1998. 5

[2] H. Bay, T. Tuytelaars, and L. Van Gool. Surf: Speeded up robust features. *Lecture Notes in Computer Science*, 3951:404, 2006. 2

[3] M. Belkin, J. Sun, and Y. Wang. Constructing Laplace operator from point clouds in \mathbb{R}^d . In *Proc. Symp. Discrete Algorithms*, pages 1031–1040, 2009. 4

[4] M. Ben-Chen, O. Weber, and C. Gotsman. Characterizing shape using conformal factors. In *Proc. Eurographics Workshop on Shape Retrieval*, 2008. 2

[5] A. Bronstein, M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust shape matching. 1

[6] A. Bronstein, M. Bronstein, M. Ovsjanikov, and L. Guibas. Shape google: a computer vision approach to invariant shape retrieval. In *Proc. NORDIA*, 2009. 2, 3, 4, 5

[7] A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel. Analysis of two-dimensional non-rigid shapes. *IJCV*, 2008. 1

[8] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006. 1

[9] D. Casasent and D. Psaltis. Position, rotation, and scale invariant optical correlation. 2

[10] O. Chum, J. Philbin, J. Sivic, M. Isard, and A. Zisserman. Total recall: Automatic query expansion with a generative feature model for object retrieval. In *Proc. ICCV*, 2007. 1, 3

[11] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21:5–30, July 2006. 1

[12] A. Elad and R. Kimmel. Bending invariant representations for surfaces. In *Proc. CVPR*, pages 168–174, 2001. 1

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

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

[15] P. Jones, M. Maggioni, and R. Schul. Manifold parametrizations by eigenfunctions of the Laplacian and heat kernels. *PNAS*, 105(6):1803, 2008. 2

[16] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Proc. SGP*, pages 156–164, 2003. 1

[17] I. Kokkinos and A. Yuille. Scale Invariance without Scale Selection. In *CVPR*, 2008. 2

[18] M. Kolomenkin, I. Shimshoni, and A. Tal. On edge detection on surfaces. In *Proc. CVPR*, 2009. 2

[19] B. Lévy. Laplace-Beltrami eigenfunctions towards an algorithm that “understands” geometry. In *Int. Conf. Shape Modeling and Applications*, 2006. 4

[20] Y. Li, H. Zha, and H. Qui. Shape topics: a compact representation and new algorithm for 3D partial shape retrieval. In *Proc. CVPR*, 2006. 2

[21] T. Lindeberg. Feature Detection with Automatic Scale Selection. *IJCV*, 30(2), 1998. 2

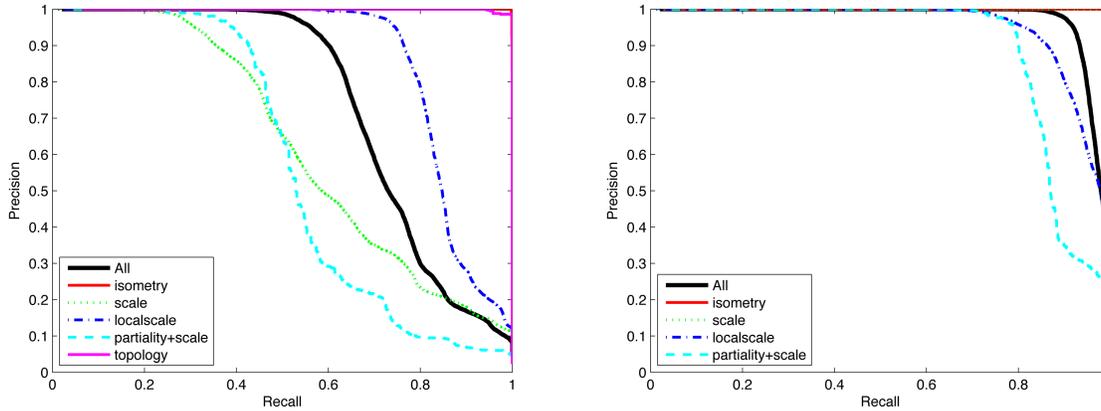


Figure 5. Shape retrieval performance using HKS (left) and SI-HKS (right).

- [22] H. Ling and D. Jacobs. Using the inner-distance for classification of articulated shapes. In *Proc. CVPR*, 2005. 1
- [23] D. Lowe. Distinctive image features from scale-invariant keypoint. *IJCV*, 2004. 2
- [24] M. Mahmoudi and G. Sapiro. Three-dimensional point cloud recognition via distributions of geometric distances. *Graphical Models*, 71(1):22–31, January 2009. 1
- [25] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing*, 22(10):761–767, 2004. 2
- [26] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5:313–346, 2005. 1
- [27] K. Mikolajczyk and C. Schmid. Scale & affine invariant interest point detectors. *IJCV*, 60(1):63–86, 2004. 1, 2
- [28] N. J. Mitra, L. Guibas, J. Giesen, and M. Pauly. Probabilistic fingerprints for shapes. In *Proc. SGP*, 2006. 2
- [29] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *ACM TOG*, 21(4):807–832, 2002. 1
- [30] 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
- [31] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, 2(1):15–36, 1993. 4
- [32] M. Reuter, S. Biasotti, D. Giorgi, G. Patanè, and M. Spagnuolo. Discrete Laplace–Beltrami operators for shape analysis and segmentation. *Computers & Graphics*, 33(3):381–390, 2009. 1, 4
- [33] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-spectra as fingerprints for shape matching. In *ACM Symp. Solid and physical modeling*, pages 101–106, 2005. 1, 5, 6
- [34] R. M. Rustomov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP*, pages 225–233, 2007. 1
- [35] J. Sivic and A. Zisserman. Video google: A text retrieval approach to object matching in videos. In *Proc. CVPR*, 2003. 1, 3
- [36] N. Sochen, R. Kimmel, and A. Bruckstein. Diffusions and confusions in signal and image processing. *JMIV*, 14(3):195–209, 2001. 1
- [37] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP*, 2009. 2
- [38] A. Tal, M. Elad, and S. Ar. Content based retrieval of VRML objects - an iterative and interactive approach. In *Proc. Eurographics Workshop on Multimedia*, 2001. 1
- [39] J. Tangelder and R. Veltkamp. A survey of content based 3D shape retrieval methods. *Multimedia Tools and Applications*, 39(3):441–471, 2008. 1
- [40] J. W. H. Tangelder and R. C. Veltkamp. A survey of content based 3D shape retrieval methods. In *Proc. Shape Modeling Applications*, pages 145–156, 2004. 1
- [41] E. Tola, V. Lepetit, and P. Fua. A fast local descriptor for dense matching. In *Proc. CVPR*, 2008. 2
- [42] R. Toldo, U. Castellani, and A. Fusiello. Visual vocabulary signature for 3D object retrieval and partial matching. In *Proc. Eurographics Workshop on 3D Object Retrieval*, 2009. 2
- [43] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun. Discrete Laplace operators: no free lunch. In *Conf. Computer Graphics and Interactive Techniques*, 2008. 4
- [44] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface Feature Detection and Description with Applications to Mesh Matching. In *Proc. CVPR*, 2009. 2
- [45] C. Zhang and T. Chen. Efficient feature extraction for 2D/3D objects in meshrepresentation. In *Proc. ICIP*, volume 3, 2001. 1
- [46] S. Zokai and G. Wolberg. Image registration using log-polar mappings for recovery of large-scale similarity and projective transformations. *Trans. Image Processing*, 14(10), 2005. 2

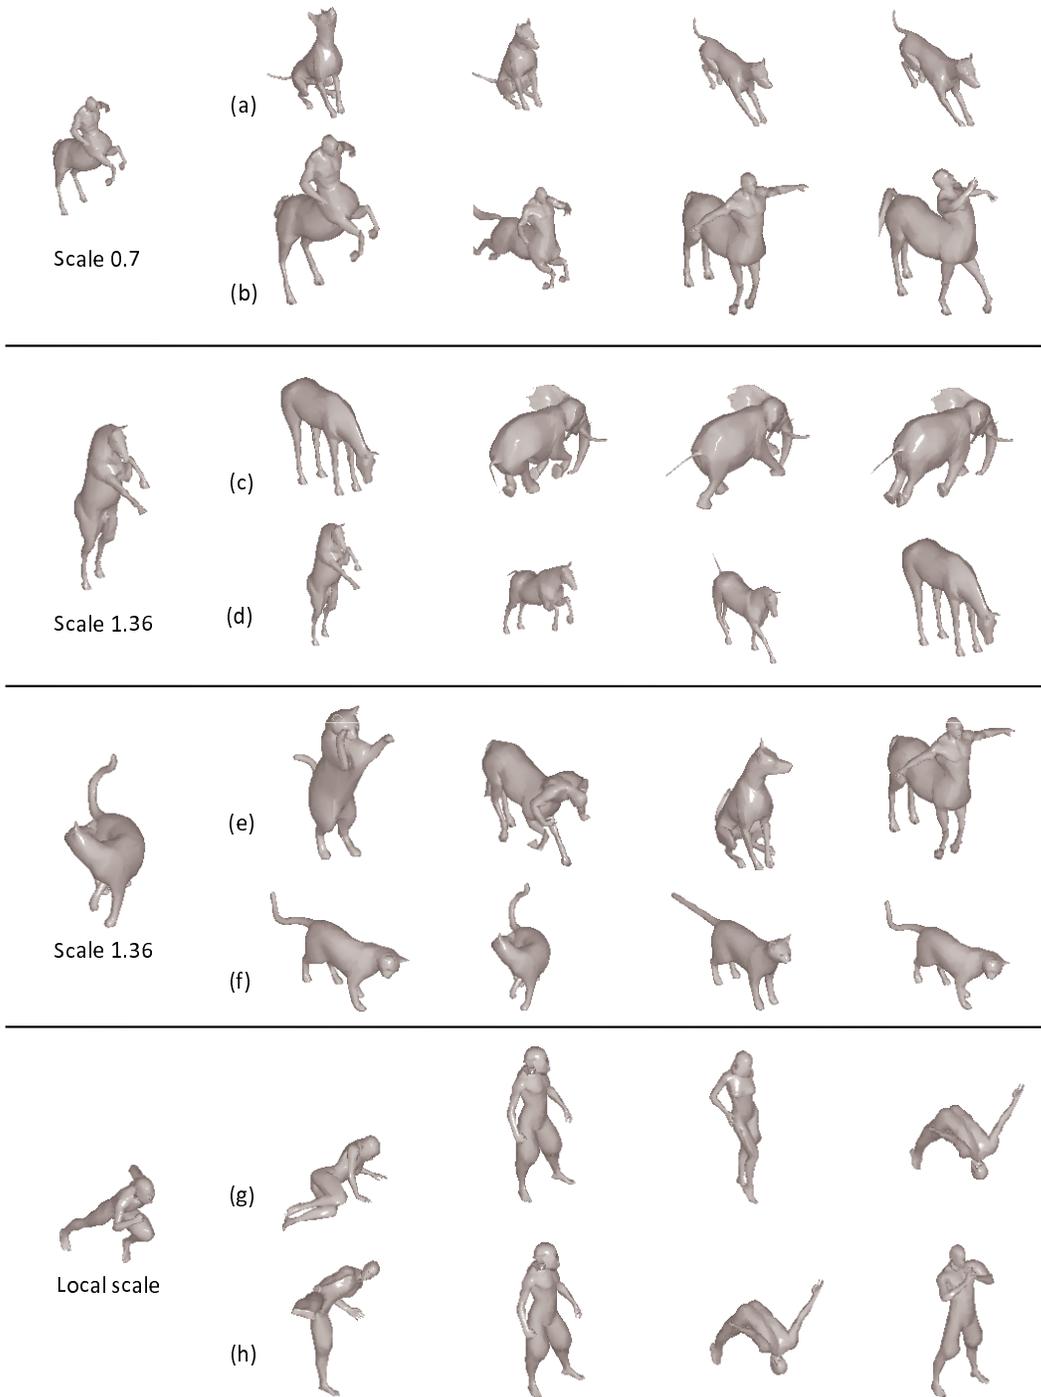


Figure 6. Shape retrieval results. Left: queries, right: first matches using HKS (a,c,e,g) and SI-HKS (b,d,f,h).