Volumetric Heat Kernel Signatures

Dan Raviv
Dept. of Computer Science,
Technion
Haifa, Israel
darav@cs.technion.ac.il

Michael M. Bronstein
Dept. of Computer Science,
Technion
Haifa, Israel
mbron@cs.technion.ac.il

Alexander M. Bronstein
Dept. of Electrical
Engineering, Tel-Aviv Univ.
Tel-Aviv, Israel
bron@eng.tau.ac.il

Ron Kimmel
Dept. of Computer Science,
Technion
Haifa, Israel
ron@cs.technion.ac.il

ABSTRACT

Invariant shape descriptors are instrumental in numerous shape analysis tasks including deformable shape comparison, registration, classification, and retrieval. Most existing constructions model a 3D shape as a two-dimensional surface describing the shape boundary, typically represented as a triangular mesh or a point cloud. Using intrinsic properties of the surface, invariant descriptors can be designed. One such example is the recently introduced heat kernel signature, based on the Laplace-Beltrami operator of the surface. In many applications, however, a volumetric shape model is more natural and convenient. Moreover, modeling shape deformations as approximate isometries of the volume of an object, rather than its boundary, better captures natural behavior of non-rigid deformations in many cases. Here, we extend the idea of heat kernel signature to robust isometry-invariant volumetric descriptors, and show their utility in shape retrieval. The proposed approach achieves state-of-the-art results on the SHREC 2010 large-scale shape retrieval benchmark.

Categories and Subject Descriptors

I.4.7 [**Image Processing and Computer Vision**]: Feature Measurement—*Feature representation, Invariants*

General Terms

Algorithms

Keywords

Volumetric Laplacian, Heat Kernel Signature

1. INTRODUCTION

Broadly, shape analysis approaches can be divided into *global* and *local* methods. The former try constructing global shape descriptors using, for example, volume and area [42], wavelets [27], statistical moments [19, 25, 38], self-similarity (symmetry) [20], distance distributions [26], Laplace-Beltrami eigenvalues [29]and eigenfunction [3, 31], metric spaces [13, 23, 7], skeletons [37] and Reeb graphs [16, 4].

Local descriptors have recently gained popularity in the shape analysis community, following the success of feature-based shape representations in computer vision and image analysis [22]. Local descriptors can be used in shape analysis in the bag of features framework [34, 9, 40], or in combination with global descriptors, such as in the recent extensions of the Gromov-Hausdorff framework [39, 12], first introduced to shape matching in [23] and [7].

There is a plethora of different local shape descriptors, including shape contexts [1], local moments [11], local diameter [14], volume descriptors [15], spherical harmonics [32], local patches [24], histograms of local geodesic distances [28] conformal factors [2], SIFT-like descriptor applied to functions defined on manifolds [41], and heat kernels [36, 10]. However, predominantly, the underlying assumption is a model of a 3D shapes as a 2D manifold, thus looking at the *boundary* surface of the physical 3D object, which is assumed to deform approximately isometrically. Here, we argue that in many natural phenomena, the isometry of the *volume* enclosed by the boundary better captures it deformations. This more restrictive set of isometries better preserve the volume of an object, and is thus more suited for many applications.

As the main instrument, we use *diffusion geometry*, which has been lately increasingly employed for shape recognition and in particular, the heat kernel descriptors proposed in [36]. For dealing with volume isometries, we replace the heat diffusion on the shape boundary by diffusion inside the shape and show that such an approach has advantages in various shape analysis applications. Philosophically similar approaches, namely considering the interior of the shape rather than its boundary, have been used in 2D shape and image analysis [21, 17, 5]. Most recently, Rustamov [30] proposed a shape descriptor based on volumetric distance distributions.

The rest of the paper is organized as follows. In Section 3, we formulate the problem and overview the heat kernel descriptors. In Section 4, we propose the volumetric heat ker-

nel descriptors, and address their numerical computation. In Section 5, we outline the construction of bag-of-feature shape descriptors from the local volumetric heat kernel descriptors. Section 6 shows experimental results, and Section 7 concludes the paper.

2. PROBLEM FORMULATION

We model a 3D object as a closed and connected three-dimensional submanifold $X \subset \mathbb{R}^3$, with boundary ∂X represented as a smooth closed two-dimensional surface. A classical and predominant approach for 3D object analysis is by studying the intrinsic geometry of the boundary surface ∂X. In these approaches, one of the desired properties of a shape descriptor is invariance to isometric deformations of the surface ∂X (referred here as *boundary isometries*), or in other words, such transformations that preserve the lengths of shortest paths between any pair of point on ∂X. These isometries are opposed to *volume isometries*, in which shortest paths between points inside X do not change. It is worthwhile noting that exact isometries of X in $\mathbb{R}^3$ are limited to the set of rigid transformations of X, that is, no non-trivial isometries of the volume exist. However, non-rigid deformations can still be modeled faithfully as approximate volume isometries. We argue that such nearly isometric deformations are more suitable to model natural non-rigid deformations and articulations of solid objects than approximate boundary isometries. This is illustrated in Figure 1 showing a glove and a human hand. A glove is an illustration of ∂X, and its deformations are boundary isometries. The hand illustrates the solid object X, and its deformations can be approximated as volume isometries. This example shows how in many cases using ∂X to model X can be wrong: some of the isometries of the glove, such as "deflating" it, are clearly inadmissible for the solid hand object. Approximate boundary isometries constitute a large class of deformations, some of which (e.g., volume-changing deformations) do not model well the natural deformations of the solid.

In this paper, we propose dealing with the object X as with the volume enclosed inside the boundary rather than the ∂X itself, and derive a descriptor based on volumetric heat propagation properties similarly to [36].

3. HEAT KERNEL SIGNATURES

Heat diffusion on the surface ∂X is governed by the *heat equation*,

$$\left(\Delta_{\partial X} + \frac{\partial}{\partial t} \right) u(t, x) = 0, \tag{1}$$

where the scalar field $u : \partial X \times [0, \infty) \to \mathbb{R}$ is the value of heat on object boundary surface ∂X at time t, and $\Delta_{\partial X}$ is the positive semi-definite *Laplace-Beltrami operator*, a generalization of the Laplacian to manifolds. The fundamental solution $k_t(x, z) : \partial X \times \partial X \times [0, \infty) \to \mathbb{R}$ of the heat equation, also called the *heat kernel*, is the solution of (1) initialized by a point heat distribution at x.

Sun *et al.* [36] proposed using the diagonal of the heat kernel $k_t(x, x)$ at multiple scales as a local descriptor, referred to as the *heat kernel signatures* (HKS). The HKS is invariant under isometric deformations of ∂X, insensitive to topological noise at small scales, and informative in the sense that under certain assumptions one could reconstruct the surface (up to an isometry) from it.

Figure 1: Deformations of a glove (left) and a solid hand (right) are an illustration of the difference between boundary and volume isometries.

Furthermore, the computation of the HKS descriptor relies on the computation of the first eigenfunctions and eigenvalues of the Laplace-Beltrami operator, which can be done efficiently and across different shape representations.

A disadvantage of the HKS is its dependence on the global scale of the shape, manifested as scaling of the descriptor and the time parameter. As a remedy, a local normalization based on logarithmic scale space and magnitude of the Fourier transform, dubbed the *scale-invariant HKS* (SI-HKS), was proposed in [10].

4. VOLUMETRIC HEAT KERNEL SIGNATURES

In order to generalize heat kernel signatures to volume, we consider the heat equation in the volume X with Neumann boundary conditions on ∂X,

$$\left(\Delta + \frac{\partial}{\partial t} \right) U(x, t) = 0 \qquad x \in \text{int}(X),$$

$$\langle \nabla U(x, t), n(x) \rangle = 0 \qquad x \in \partial X \tag{2}$$

where n is the normal to the boundary surface ∂X, Δ is the standard Laplacian in $\mathbb{R}^3$, and $U : (X \subset \mathbb{R}^3) \times [0, \infty) \to \mathbb{R}$ is the volumetric heat distribution in X.

We denote by $K_t(x, y) : X \times X \times [0, \infty) \to \mathbb{R}$ the heat kernel of (2). By the spectral decomposition theorem, $K_t(x, y)$ can be written as [18]

$$K_t(x, y) = \sum_{l=0}^{\infty} e^{-\lambda_l t} \Phi_l(x) \Phi_l(y). \tag{3}$$

where λ_l, Φ_l are the eigenvalues and eigenfunctions of the Laplacian operator with the above boundary conditions,

$$\Delta \Phi_l(x) = \lambda_l \Phi_l(x);$$

$$\langle \nabla \Phi_l(x), n(x) \rangle = 0 \qquad x \in \partial X \tag{4}$$

The heat kernel of (2) gives rise to volumetric *heat kernel*

signatures (VHKS) which we define exactly in the same way as HKS. Namely, for each point $x \in X$,

$$\mathbf{h}(x) = (K_{t_1}(x, x), \dots, K_{t_n}(x, x)), \qquad (5)$$

is used as an n-dimensional local descriptor, where $t_1, \dots, t_n$ are different time-scales.

VHKS are related to a geometric quantity called the *scalar curvature* or the trace of the Ricci curvature by

$$K_t(x, x) = \frac{1}{(4\pi t)^{3/2}} \left(1 + \frac{1}{6} s(x) \right), \qquad (6)$$

where $s(x)$ is the scalar curvature. $s(x)$ can be defined as the ratio of the volume of a small ball centered at x to the volume of a three-dimensional Euclidean ball of the same radius.

4.1 Numerical computation

In our experiments, we sampled the interior of the volumes on a regular Cartesian grid with square voxels, which allows to use the standard Laplacian in $\mathbb{R}^3$ as the Laplace-Beltrami operator. We use the finite difference scheme to evaluate the second derivative in each direction in the volume, and used the shadow variables technique to enforce Neumann boundary conditions.

Since the coefficients $e^{-\lambda_l t}$ decay fast for $t > 0$, we can approximate the VHKS descriptor (3) by the truncated sum

$$K_t(x, y) \approx \sum_{l=0}^{k} e^{-\lambda_l t} \Phi_l(x) \Phi_l(y). \qquad (7)$$

Thus, in practice, the computation of the VHKS descriptor boils down to computing k smallest eigenvalues and corresponding eigenvectors of the Laplacian operator in $\mathbb{R}^3$ restricted to the volume X. In our experiments, $k = 150$ was used.

5. VOLUMETRIC BAGS-OF-FEATURES

In order to aggregate local point descriptors into a single global volume descriptor, we follow the ShapeGoogle framework introduced in [9]. Given a volume X, a dense point descriptor $\mathbf{h}(x) \in \mathbb{R}^n$ is first computed for every $x \in X$. The second step consists of vector quantization of the descriptors. Given a vocabulary $\mathcal{V} = \{\mathbf{h}_1, \dots, \mathbf{h}_V\}$ of representative descriptors, $\mathbf{h}(x)$ at each point is replaced by the V-dimensional distribution

$$\boldsymbol{\theta}_k(x) = \exp\{-\|\mathbf{h}_k - \mathbf{h}(x)\|^2 / 2\sigma^2\}, \qquad (8)$$

where σ is a parameter. In the limit $\sigma \to 0$, the process boils down the the standard hard vector quantization assigning each descriptor the index of its nearest neighbor in the vocabulary. The vocabulary is created offline by clustering a training set of descriptors.

After vector quantization, the point-wise distributions $\boldsymbol{\theta}(x)$ are integrated over the entire volume, resulting in the volumetric bag-of-features

$$\mathbf{f}(x) = \int \boldsymbol{\theta}(x) dx. \qquad (9)$$

Volumetric bags-of-features are compared using a standard metric in $\mathbb{R}^V$, e.g. the L_1 metric.
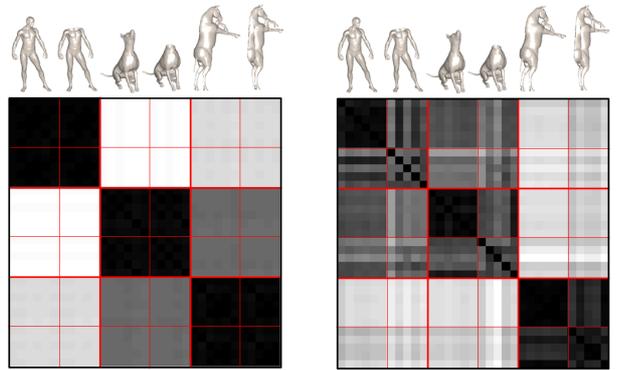

Figure 2: Similarity matrix using HKS (left) and VHKS (right) on a dataset containing three shape classes (human, dog, horse) undergoing boundary and volumetric isometric deformations. Bright colors stand for large dissimilarity. HKS are unable to distinguish between deformations that substantially violate the natural "state of aggregation" of a shape.

6. RESULTS

In order to test our approach, we used local feature descriptors to construct global shape descriptors, which were used to compare shapes. Global shape descriptors were constructed using bags of geometric words proposed in [9]. A geometric vocabulary of size 64 was built using clustering in the space of descriptors. We used the first 150 eigenpairs of the Laplacian to approximate the heat kernel and 6 time scales ($t = 1024, 1351, 1783, 2353, 3104$ and 4096) in the construction of the VHKS. Shapes were represented as volumes; the VHKS at each point in the volume of the shape were replaced by the closest geometric word from the vocabulary using soft vector quantization. The distribution of geometric words (64-dimensional bag of features) was used as the shape descriptor. L_1 distance was used to compare the bags of features.

6.1 Boundary vs volume isometries

In the first experiment, we designed a special dataset to illustrate the difference between boundary and volume isometries and thus emphasize the problem of approaches considering 3D shapes as 2D boundary surfaces. Our dataset contained objects from three classes (humans, dogs, and horses). In each class, the null pose of the object underwent five approximate volume isometries (being also boundary isometries) and five boundary isometries which change the volume significantly (such as inverting the human head inside the torso). For all the shapes we computed local HKS descriptors on the boundary following [9], and volumetric HKS descriptors proposed here. These local descriptors were aggregated into a global bag of features as described above.

Figure 2 shows the distance matrices between the shapes using HKS (left) and VHKS (right). HKS fail to distinguish between volume-changing deformation, though they render the shapes visually very different. With VHKS, on the other hand, such deformations are distant from the volume-preserving ones. This fact is further visible from Figure 3 depicting the two descriptors under volume-changing and volume-preserving isometries of the boundary. Observe that

Figure 3: RGB representation of the first three components of the HKS descriptor (first row) and VHKS descriptor on the shape boundary (second row). Second and third columns show, respectively, a volume-preserving and a volume changing boundary isometries.

while both descriptors are capable of discriminating between non-isometric shapes, HKS are invariant under volume-changing isometries of the boundary, while VHKS are not.

6.2 Shape retrieval

In order to evaluate the performance of the proposed descriptors for shape retrieval, we used the SHREC 2010 robust large-scale shape retrieval benchmark, simulating a retrieval scenario, in which the queries include multiple modifications and transformations of the same shape [6].

Dataset. The dataset used in this benchmark was aggregated from three public domain collections: TOSCA shapes [8], Robert Sumner's collection of shapes [35], and Princeton shape repository [33]. The shapes in the original SHREC dataset are given as triangular meshes with the number of vertices ranging approximately between 300 and 30,000. We converted the shapes into a volumetric representation using rasterization, with fixed voxel size. Typical volume size was $150 \times 150 \times 150$ The original dataset consisted of two parts: 715 shapes from 13 shape classes with simulated transformation (55 per shape) used as queries and the rest of 456 shapes, totalling in 1184 shapes.

Queries. The query set consisted of 13 shapes taken from the dataset (null shapes), with simulated transformations applied to them. An insignificant number of shapes containing thin structures was removed due to our inability to accurately convert them into a volumetric representation. For each null shape, transformations were split into 10 classes shown in Figure 4 (we removed the partial view transformation present in the original benchmark, as it represents a partial occlusion of the boundary and not what is expected to be a partial occlusion of the volume). In each class, the transformation appeared in five different versions numbered 1–5. The total number of transformations per shape was 50, and the total query set size was 600. Each query had one correct corresponding null shape in the dataset.

Transformation	Strength				
	1	≤2	≤3	≤4	≤5
Isometry	100.00	100.00	100.00	100.00	100.00
Topology	100.00	100.00	100.00	100.00	100.00
Holes	100.00	100.00	100.00	100.00	98.75
Micro holes	100.00	100.00	100.00	100.00	100.00
Scale	0.61	11.94	8.81	6.74	5.46
Local scale	100.00	93.35	81.86	69.04	60.81
Sampling	100.00	100.00	100.00	100.00	100.00
Noise	100.00	100.00	100.00	100.00	100.00
Shot noise	100.00	100.00	100.00	100.00	100.00
Mixed	100.00	61.14	41.65	31.47	25.29

Table 1: Performance (mAP in %) of our approach using VHKS descriptors.

Evaluation criteria. Evaluation simulated matching of transformed shapes to a database containing untransformed (null) shapes. As the database, all 469 shapes with null transformations were used. Multiple query sets according to transformation class and strength were used. For transformation x and strength n, the query set contained all the shapes with transformation x and strength $\leq n$.

Performance was evaluated using precision/recall characteristic. *Precision* $P(r)$ is defined as the percentage of relevant shapes in the first r top-ranked retrieved shapes. In the present benchmark, a single relevant shape existed in the database for each query. *Mean average precision* (mAP), defined as

$$mAP = \sum_r P(r) \cdot rel(r),$$

(where $rel(r)$ is the relevance of a given rank), was used as a single measure of performance. Ideal performance retrieval performance results in first relevant match with mAP=100%.

Comparison to other methods. Each shape in the benchmark was represented as a bag of VHKS features in a vocabulary of size 64, as described above. Table 1 show the performance (mAP in %) across transformation classes and strengths of shape retrieval using bags of features based on VHKS local descriptors. For comparison, in Table 3 we show the performance of ShapeGoogle [9] using local HKS descriptors computed on a mesh using cotangent weigh approximation of the Laplace-Beltrami operator. One can observe that our approach is more robust under strong isometric and topological transformations and is less sensitive to noise and resampling. Since the descriptor is not scale-invariant, the performance on global and local scaling and mixed transformations is poor.

Table 2 shows the performance of volumetric ShapeDNA [29], which represents the shapes as vectors of the corresponding Laplacian eigenvalues. We used the first 150 eigenvalues. Our approach significantly outperforms this method.

The complexity of calculating VHKS is higher than that of HKS. In surfaces, several thousand points are enough for constructing a distinguishable signature, while in volumes that number increases by a factor of ten. Since the Laplacian is represented by a sparse matrix, calculating several dozen eigenvalues and eigenvectors remains a fast procedure. Under core 2 Duo 3GHz computer, it takes approximately several seconds to calculate each VHKS. Training the dictionary is done off-line and takes several minutes.
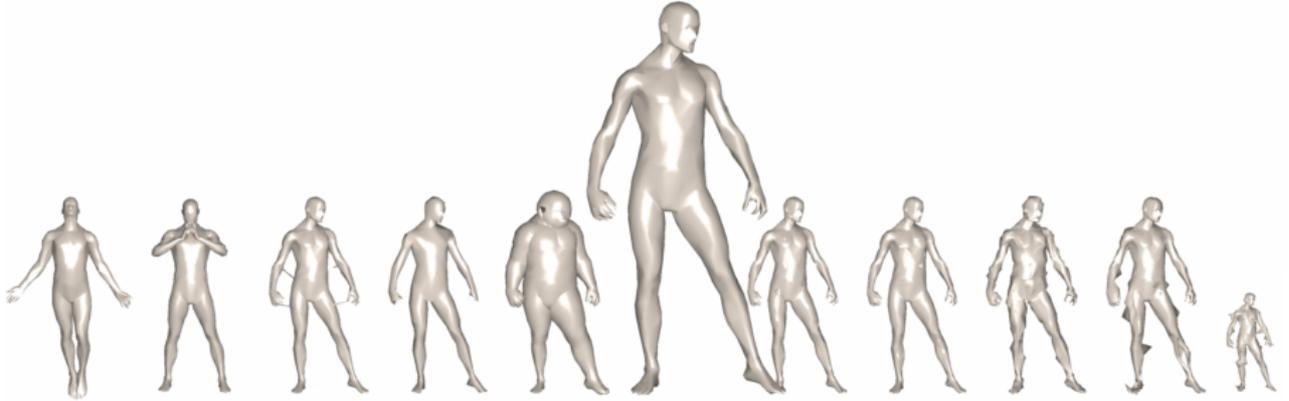
Figure 4: Transformations of the human shape used as queries (shown in strength 5, left to right): null, isometry, topology, sampling, local scale, scale, holes, micro holes, noise, shot noise, and mixed.

	Strength				
Transformation	1	≤2	≤3	≤4	≤5
Isometry	56.97	53.65	49.82	47.97	48.91
Topology	61.97	59.76	58.83	56.92	55.44
Holes	53.89	53.72	51.91	45.40	41.41
Micro holes	51.20	55.05	56.44	56.15	56.06
Scale	0.91	1.93	1.52	1.24	1.08
Local scale	49.69	37.11	33.63	28.15	23.41
Sampling	58.26	60.18	59.45	60.70	58.84
Noise	50.74	51.18	48.76	43.17	38.50
Shot noise	58.93	57.04	50.53	42.54	36.88
Mixed	55.25	30.09	21.00	16.23	13.26

Table 2: Performance (mAP in %) of volumetric ShapeDNA.

	Strength				
Transformation	1	≤2	≤3	≤4	≤5
Isometry	100.00	100.00	98.61	98.96	99.17
Topology	100.00	96.04	94.67	90.93	89.32
Holes	100.00	96.43	94.91	85.81	77.41
Micro holes	100.00	100.00	100.00	100.00	100.00
Scale	80.03	90.02	93.34	95.01	96.01
Local scale	100.00	100.00	95.90	86.43	78.38
Sampling	100.00	100.00	100.00	100.00	95.92
Noise	100.00	100.00	100.00	100.00	100.00
Shot noise	100.00	100.00	100.00	100.00	100.00
Mixed	40.84	39.95	46.38	47.79	44.79

Table 3: Performance (mAP in %) of the 2D HKS descriptor computed with cotangent weight discretization of the Laplacian.

7. CONCLUSIONS

Our starting point was the fact that not all boundary isometries of solids correspond to meaningful deformations. We argued that volume isometries better model natural non-rigid deformations and constructed local descriptors based on volumetric heat kernels. Such descriptors readily lend themselves to the bag-of-features representation used in shape retrieval. Experimental results on the SHREC 2010 benchmark demonstrate that the proposed descriptors are more robust to topological and geometric noise than previously proposed approaches, achieving state-of-the-art results in several deformation categories. In future studies, we intend to construct a scale-invariant version of the VHKS descriptors.

8. ACKNOWLEDGEMENTS

9. REFERENCES

[1] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *Trans. PAMI*, pages 509–522, 2002.

[2] M. Ben-Chen, O. Weber, and C. Gotsman. Characterizing shape using conformal factors. In *Proc. 3DOR*, 2008.

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

[4] S. Biasotti, S. Marini, M. Mortara, and G. Patané. An overview on properties and efficacy of topological skeletons in shape modeling. In *Proc. SMI*, pages 245–254, 2003.

[5] A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel. Analysis of two-dimensional non-rigid shapes. *IJCV*, 78(1):67–88, 2008.

[6] A. M. Bronstein, M. M. Bronstein, U. Castellani, B. Falcidieno, A. Fusiello, A. Godil, L. J. Guibas, I. Kokkinos, Z. Lian, M. Ovsjanikov, G. Patané, M. Spagnuolo, and R. Toldo. Shrec 2010: robust large-scale shape retrieval benchmark. In *Proc. 3DOR*, 2010.

[7] 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.

[8] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer-Verlag New York Inc, 2008.

[9] A. M. Bronstein, M. M. Bronstein, M. Ovsjanikov, and L. J. Guibas. Shape google: a computer vision approach to invariant shape retrieval. In *Proc. NORDIA*, 2009.

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

[11] U. Clarenz, M. Rumpf, and A. Telea. Robust feature detection and local classification for surfaces based on moment analysis. *Trans. Visualization and Computer Graphics*, 10(5):516–524, 2004.

[12] A. Dubrovina and R. Kimmel. Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In *Proc. 3DPVT*, 2010.

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

[14] R. Gal, A. Shamir, and D. Cohen-Or. Pose-oblivious shape signature. *Trans. Visualization and Computer Graphics*, 13(2):261–271, 2007.

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

[16] M. Hilaga, Y. Shinagawa, T. Kohmura, and T. Kunii. Topology matching for fully automatic similarity estimation of 3D shapes. In *Proc. Computer Graphics and Interactive Techniques*, pages 203–212, 2001.

[17] A. Ion, G. Peyré, Y. Haxhimusa, S. Peltier, W. G. Kropatsch, and L. Cohen. Shape matching using the geodesic eccentricity transform–a study. In *Proc. Workshop of the Austrian Association for Pattern Recognition (OAGM/AAPR)*, 2007.

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

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

[20] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Symmetry descriptors and 3D shape matching. In *Proc. SGP*, pages 115–123, 2004.

[21] H. Ling and D. W. Jacobs. Shape classification using the inner-distance. *Trans. PAMI*, 29(2):286–299, 2007.

[22] D. Lowe. Distinctive image features from scale-invariant keypoint. *IJCV*, 2004.

[23] 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.

[24] N. J. Mitra, L. J. Guibas, J. Giesen, and M. Pauly. Probabilistic fingerprints for shapes. In *Proc. SGP*, 2006.

[25] M. Novotni and R. Klein. 3D Zernike descriptors for content based shape retrieval. In *Proc. ACM Symp. Solid Modeling and Applications*, pages 216–225, 2003.

[26] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *TOG*, 21(4):807–832, 2002.

[27] E. Paquet, M. Rioux, A. Murching, T. Naveen, and A. Tabatabai. Description of shape information for 2-D and 3-D objects. *Signal Processing: Image Communication*, 16(1-2):103–122, 2000.

[28] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Symmetries of non-rigid shapes. In *Proc. NRTL*, 2007.

[29] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-spectra as fingerprints for shape matching. In *Proc. ACM Symp. Solid and Physical Modeling*, pages 101–106, 2005.

[30] R. Rustamov. Robust volumetric shape descriptor. In *Proc. 3DOR*, 2010.

[31] R. M. Rustamov. Laplace-Beltrami eigenfunctions for

[32] P. Shilane and T. Funkhauser. Selecting distinctive 3D shape descriptors for similarity retrieval. In *Proc. Shape Modelling and Applications*, 2006.

[33] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. The Princeton shape benchmark. In *Proc. SMI*, pages 167–178, 2004.

[34] J. Sivic and A. Zisserman. Video google: A text retrieval approach to object matching in videos. In *Proc. CVPR*, 2003.

[35] R. Sumner and J. Popović. Deformation transfer for triangle meshes. In *Proc. Conf. Computer Graphics and Interactive Techniques*, pages 399–405, 2004.

[36] J. Sun, M. Ovsjanikov, and L. J. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP*, 2009.

[37] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. In *Proc. SMI*, volume 130, 2003.

[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.

[39] N. Thorstensen and R. Keriven. Non-rigid shape matching using geometry and photometry. In *Proc. CVPR*, 2009.

[40] R. Toldo, U. Castellani, and A. Fusiello. Visual vocabulary signature for 3D object retrieval and partial matching. In *Proc. 3DOR*, 2009.

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

[42] C. Zhang and T. Chen. Efficient feature extraction for 2D/3D objects in meshrepresentation. In *Proc. ICIP*, volume 3, 2001.

deformation invariant shape representation. In *Proc. SGP*, pages 225–233, 2007.