

Local Signature Quantization by Sparse Coding

D. Boscaini¹ and U. Castellani¹

¹Department of Computer Science, University of Verona

Abstract

In 3D object retrieval it is very important to define reliable shape descriptors, which compactly characterize geometric properties of the underlying surface. To this aim two main approaches are considered: global, and local ones. Global approaches are effective in describing the whole object, while local ones are more suitable to characterize small parts of the shape. Some strategies to combine these two approaches have been proposed recently but still no consolidate work is available in this field. With this paper we address this problem and propose a new method based on sparse coding techniques. A set of local shape descriptors are collected from the shape. Then a dictionary is trained as generative model. In this fashion the dictionary is used as global shape descriptor for shape retrieval purposes. Preliminary experiments are performed on a standard dataset by showing a drastic improvement of the proposed method in comparison with well known local-to-global and global approaches.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—

1. Introduction

The recent improvements of 3D acquisition systems and the large proliferation of 3D models increase the interest in 3D shape retrieval methods [IJL*05, FKMS05, TV04, LGF*10, LGB*11]. A challenging issue is to elaborate a suitable canonical representation of the objects to be indexed. In the literature, this characterization is referred to as *descriptor* or *signature*. In general, descriptors are *global* or *local*. The former consist in a set of features that effectively and compactly describe the whole 3D model [FMK*03]. The latter, instead, are collections of local features of relevant object subparts (i.e. single points or regions) [SF06, CCM11].

In this paper we address the problem of defining a global shape descriptor starting from a set of local point signatures [SOG09, BK10, ASC11, CCM11]. The overall aim is to obtain the advantages of the two approaches. From one side we are able to compare global shapes rather than a set of single points. From the other side, we exploit local informations which, in general, is more robust to noise and missing parts and more suitable to deal with partial objects. To this aim a popular method is to introduce a sort of counting approach by collecting local informations into a histogram which leads to a local-to-global signature. Such examples are distance shape distributions [FKMS05] or the *bag of words* approach [BBGO11, TCF10, DK12].

With this work we propose to go beyond the bag of words approach by exploiting recently proposed dictionary learning methods employing sparse coding techniques [MBPS09, MBPS10]. Starting from a set of local signatures we learn a *dictionary* which is able to summarize the most relevant properties of such set. This leads to a more sparse representation of the shape which is used for its description. We propose this approach in a shape retrieval context. The idea is to train a dictionary for each class of shape. Then, in the query phase, a given shape is generated by all available dictionaries and it is assigned to the class with less generative error.

A well define shape retrieval pipeline is proposed by combining effectively the most promising local shape descriptors with the proposed *local-to-global* approach based on sparse coding. The main steps are:

- Local descriptor computation by diffusion geometry signatures [ASC11],
- Dictionary learning by sparse coding [MBPS09, MBPS10],
- Shape matching by best generative signature estimation.

In particular, as local descriptors is employed the recently proposed Wave Kernel Signature [ASC11] which already shown its effectiveness for point-to-point matching. Then, dictionary learning method is applied by using the *Lasso*

approach [Tib96]. Although such approach is quite popular in signal processing, to the best of our knowledge, only recently it has been proposed in computer vision for 2D image coding and very few work has been done in 3D domain (e.g. [PBB*13]).

Experiments are evaluated on standard dataset, i.e. SHREC 2011 robustness benchmark [BBB*11] and we show a drastic improvement over both good global shape descriptor, such as Shape DNA [RWP06], and standard local-to-global approaches, such as shape distance distribution [FKMS05].

The rest of the paper is organized as follows. Section 2 reports the state of the art by focusing mainly on the shape descriptors which are related to the proposed work. Section 3 describes the background for both sparse coding and local descriptors, and it introduces the proposed local-to-global approach. Section 4 reports the experimental results by showing the performance of the proposed descriptor in comparison with other methods. Finally, in Section 5, conclusions are drawn and future work is envisaged.

2. Related work

Ideal shape descriptors should satisfy some properties such as *discriminateness*, *robustness* to noise, *invariance* to isometries and other shape transformations, *compactness* and so on (see, e.g. [RWP06, Bro]). In the following we briefly revise global, local, and local-to-global approaches.

2.1. Global approaches

Regarding *global* approaches, spectral methods are largely employed [FKMS05, TV04]. For instance Shape DNA [RWP06] computes the spectral decomposition of the Laplace-Beltrami operator defined on the manifold represented by the shape and uses the truncated set of the computed eigenvalues as global signature. This leads to a very effective descriptor which was successfully employed on several applicative scenarios such as shape retrieval and shape matching in medical domain [RWP06, LGB*11].

2.2. Local approaches

Local descriptors are often employed for point-to-point correspondences [FKMS05, TV04, CCM11]. A common approach is to collect local geometric properties on the point neighborhood and accumulate these values on a multidimensional histogram. Examples are *Spin Images* [JH99] or *Shape Context* [BMP02, FHK*04, KCB09]. Other approaches exploit probabilistic properties of the shape, e.g. in [CCM11] Hidden Markov Models are adapted to work on 3D surfaces. Another very popular class of local descriptors are based on *diffusion geometry* [SOG09, BK10, ASC11]. In [SOG09] the so called Heat Kernel Signature (HKS) was introduced which exploits the local surface properties at different scales. Some extensions of HKS are proposed in [BK10]

to deal with scale invariance. Recently, in [ASC11], was proposed the so called Wave Kernel Signature (WKS). It employs a different physical model being related with oscillation rather than diffusion processes.

2.3. Local-to-global approaches

Local-to-global approaches are therefore introduced to define a global signature from a collection of local descriptors. A simple method consists of computing pairwise distances among points in the descriptor space and accumulate these distances into a histogram [FKMS05, BBM*10]. More sophisticated techniques exploit probabilistic methods, such as in [MGGP06], where a probabilistic fingerprint is introduced. Being encouraged by feature-based methods developed in Computer Vision, several work employed the so called bag of words paradigm [BBGO11, TCF10, DK12]. In [BBGO11] the bag of word descriptor is computed from the set of local HKS signatures. In [TCF10] a region-based approach is introduced where the *visual words* are defined by region properties computed after shape segmentation. In [DK12] authors extract the bag of words signature after detection of feature points and by collection properties that make the descriptor scale invariant.

In this paper we propose to exploit dictionary learning and sparse coding approaches [MBPS09, MBPS10]. To the best of our knowledge, such approach is very few adopted for 3D shapes and only recently some methods have been proposed, such as in [PBB*13], for point-to-point correspondences of non-rigid or partial shapes. Here, we propose sparse coding method for local-to-global shape description.

3. Global shape descriptor by sparse coding

In this section we introduce: i) the general theoretical background of sparse coding, ii) local descriptors involved in our method, and iii) the main contribution of the paper, i.e. how to exploit the former approach to propose a global signature from a set of local descriptors.

3.1. Background

When a machine learning approach is employed a general issue to be addressed is the following [SS02]: given two classes of objects x_i , and a new object x , how can we assign the unknown object to the right class? To distinguish the objects belonging to a class from the others, we assign a *label* y_i to each object, i.e.

$$(x_1, y_1), \dots, (x_n, y_n) \in X \times \{\pm 1\},$$

where the labels are chosen as +1 and -1 for the sake of simplicity and X is some non-empty set containing the *patterns* x_i . Given some new pattern $x \in X$, we want to infer the corresponding label $y \in \{\pm 1\}$.

To this end an interpolation on the given data, i.e.

$$\min_f \|y_i - f(x_i)\|_2^2$$

is useless since it is not able to generalize well for unseen patterns. A possible approach at this problem is suggested by Tychonoff *regularization* theory and consist of a restriction of the class of admissible solutions, i.e., a compact set.

Indeed, the previous problem can be reformulated as:

$$\min_f \|y_i - f(x_i)\|_2^2 + \lambda R(f),$$

where

- $\|y_i - f(x_i)\|_2^2$ is the data term,
- $R(f)$ is the *regularization* term,
- $\lambda > 0$ is the so-called *regularization* parameter, which specifies the trade-off between fidelity to the data in the sense of ℓ^2 norm, as represented by the former term, and simplicity of the solution, enforced by $R(f)$.

An example of regularization operator is $R(f) = \|f^{(m)}\|_2$, for some $m \in \mathbb{N}$. This particular choice promotes the smoothness of the solution.

Let us now address a slightly different problem. Now suppose we have a sentence s and a dictionary D . We want to explain the sentence s with words contained in D . This problem could be formalized as

$$\min_{\alpha} \|s - \alpha D\|_2^2,$$

the idea is that vector α picks up only the words that describe the sentence s . In general, a dictionary is *over complete*: there are a lot of words with the same or similar meaning. For this reason we might be interested to consider the minimum number of words as possible. Regularization theory help us also in this case. Indeed, if we consider the following problem:

$$\min_{\alpha} \|s - \alpha D\|_2^2 + \lambda R(\alpha), \quad (1)$$

by choosing $R(\alpha) = \|\alpha\|_1$, we are promoting the sparsity of the solution. In this case we refer as *sparse coding* and the corresponding problem is known as *Lasso formulation* [Tib96].

In general, as described in [PBB*13], s could be thought as a generic signal and the interpretation of Lasso formulation could be the following: many families of signals can be represented as a sparse linear combination in an appropriate domain, usually referred to as *dictionary*, so that $s \approx \alpha D$. In other words, the signal s could be *generated* by αD . Finally, given the signal s and the dictionary D , the solution of the unconstrained convex minimization problem of equation (1) gives us the sparse vector α .

However in general the dictionary D is not available. We therefore are interested in inferring both the vector α and the

dictionary D from the signal s . The problem becomes:

$$\min_D \left(\min_{\alpha} \|s - \alpha D\|_2^2 + \lambda \|\alpha\|_1 \right), \quad (2)$$

In [MBPS09, MBPS10], problem (2) was solved employing an alternating minimization method between the variables D and α^i .

As a further step we should consider that in the more general case, instead of a single signal s , we have a collection of signals $\mathbf{s} = \{s^i\}_{i=1, \dots, N}$. Therefore, Equation (2) can be generalized as:

$$\min_D \frac{1}{N} \sum_{i=1}^N \min_{\alpha^i} \left(\|s^i - \alpha^i D\|_2^2 + \lambda \|\alpha^i\|_1 \right), \quad (3)$$

where $\alpha = \{\alpha^i\}_{i=1, \dots, N}$ is thought as a collection of vectors.

3.2. Local shape descriptors

Since we want to employ sparse coding technique for shape analysis, we need to extract a signal from the underlying geometry of a shape that possibly be robust to non rigid deformations. To this aim we consider as the signal a collection of local descriptors collected at each vertex of the considered shape. In order to satisfy all the previous hypothesis, properties of Laplace-Beltrami operator on the 2-manifold represented by the shape are exploited.

In the context of diffusion geometry, the most popular local descriptor is Heat Kernel Signature (HKS) [SOG09] and its scale invariant version, SI-HKS [BK10]. They are based on the properties of the heat diffusion process on the shape governed by the heat equation.

$$\left(\frac{\partial}{\partial t} - \Delta \right) u(x, t) = 0. \quad (4)$$

For the signal processing perspective, we could say that the solution $u(x, t)$ of differential equation (4) at time t can be expressed by the convolution of the *impulse response* $h_t(x, t)$ by the initial data $u_0(x)$,

$$u(x, t) = \int h_t(x, y) u_0(y) d\sigma(y).$$

The kernel of this integral operator is called *heat kernel* and it corresponds to the amount of heat transferred from point x to point y after time t . In particular, HKS represents the autodiffusion process $h_t(x, x)$ centered in a vertex x of the shape, at different time scales.

As described in [Bro], the heat kernel descriptor could be thought as a collection of low-pass filters. This emphasis of low frequencies damages the ability of the descriptor to precisely localize shape features. A remedy to the poor feature localization of the heat kernel descriptor was proposed by the so called *Wave Kernel Signature* (WKS) in [ASC11]. The authors proposed to replace the heat diffusion equation

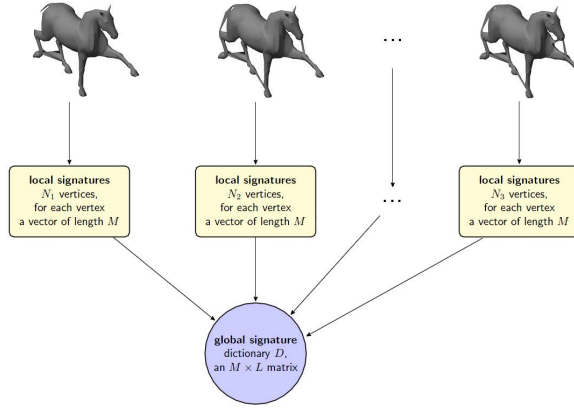


Figure 1: The dictionary D^c representing the class c of an horse, could be learned from the collection of local signatures of an entire class of topological deformations of the horse null shape. Here is represented the pipeline of the proposed method: from each deformed shape of N_i vertices we extract N_i WKS vectors of M components. Then we collected all this vectors in a single matrix \mathbf{s} and learn a dictionary D^c of L “words” that represent the class c up to all topological deformations considered.

(4), by the Shrödinger equation

$$\left(\frac{\partial}{\partial t} + i\Delta\right)v(x, t) = 0,$$

where $v(x, t)$ is the complex wave equation. Changing only slightly the differential equation that govern the process, now the physical interpretation is very different: it represents the average probability of measuring a quantum particle with a certain energy distribution at a specific location. That is, instead of representing diffusion, v has oscillatory behavior.

Letting vary the energy of the particle, the WKS encodes and separates information from various different frequencies. In terms of the former interpretation of HKS as a collection of filters, in [Bro] was noted that WKS can be thought as a collection of *band-pass* filters. As a result, the wave kernel descriptor exhibits superior feature localization. For this reason, we chose to consider this local signature for our experiments.

3.3. Local-to-global descriptors

Once WKS descriptors are computed, we consider them as the collection of signals \mathbf{s} . Then, learning techniques described in [MBPS09, MBPS10] are employed for solving problem (3), i.e. learn the dictionary D . Here we consider D as the global signature for each shape \mathbf{O} . In this way we have a matrix as a global signature of each shape \mathbf{O}_i of a 3D shapes database. The query phase could be done employing the classical leave-one-out approach. That is, we could compare the descriptor of the query shape with the descriptors of all the other shapes in the database and assign to the query shape the class of the shape whose descriptor generates the smallest error.

Here we follow another way. It is worth noting that \mathbf{s} could be considered also as the collection of local signatures of an entire class of shapes subjected to a non-rigid deformation, as shown in Figure 1. In this case, several deformations of the same object can contribute in learning the dictionary. More in details:

- $\{\mathbf{O}_1^c, \dots, \mathbf{O}_k^c\}$ are several instances of class c ,
- k is the number of instances of the same class deformations,
- D^c represents the dictionary of class c ,

D^c is trained by all signatures of the instances of the class $\{\mathbf{O}_1^c, \dots, \mathbf{O}_k^c\}$. If c represents a class of noise deformations of the same shape, then \mathbf{O}_i^c represents a noise deformation of \mathbf{O} and k the number of noise deformations of the shape \mathbf{O} present in c .

In order to deal with multiple classes of shapes (e.g. men, cats, dogs, and so on) several dictionaries can be trained $\{D^{c_1}, D^{c_2}, \dots\}$, one for each available class. Then, in the query phase, given a shape \mathbf{O} and its collection of local signatures \mathbf{s} , we solve problem (1) for each dictionary D^{c_i} and we obtain the vectors α^{c_i} . At the shape \mathbf{O} is then assigned the class c such that

$$\|\mathbf{s} - \alpha^c D^c\| = \min_i \|\mathbf{s} - \alpha^{c_i} D^{c_i}\|. \quad (5)$$

In retrieval applications, the principal advantage of this method is that it allow to compare the signature of a query shape only with the dictionary of the classes of the shapes present in the database considered. Conversely, in the majority of the existing approaches, the matching is done between the query shape and all the instances of the database. Another important advantage is that the dictionary encodes

more instances of the same non-rigid deformation, making the proposed signature a descriptor of the entire class of deformations rather than a single shape descriptor.

4. Results

The reported experiments are evaluated on the SHREC 2011 robustness benchmark. In particular the database is composed of 12 different triangulated meshes from TOSCA [BBK08] and Sumner [SP04] databases, that we consider as null shapes, and their non-rigid deformations. For each null shape reported in Figure 2, transformations were split into 9 different types:

- *affine*,
- *big holes*,
- *micro holes*,
- *scale*,
- *down sampling* (less than 20% of original points),
- *additive Gaussian noise*,
- *shot noise*,
- *topology* (welding of shapes vertices resulting in different triangulation) and
- *view*,

as reported in Figure 3. Each triangular meshes has about 1500 vertices.

Each type of transformation appeared in five different versions numbered from 1 to 5. In all deformation types, the version number correspond to the transformation strength level: the higher the number, the stronger the deformation (e.g. in noise transformation, the noise variance is proportional to the strength number). For scale transformation, the levels 1 – 5 correspond to scaling by the factor 0.5, 0.875, 1.25, 1.625 and 2.

For each class of deformations, we have 60 shapes, 5 for every null class. The entire database contains 552 shapes: the 540 deformed shapes and the 12 null shapes.

According to the pipeline proposed in section 3, we extract the WKS [ASC11] signature for each vertex of each shapes. For the signature extraction we based on the Matlab code freely available on <http://vision.in.tum.de/publications>. In all our experiments the parameters were fixed accordingly with [ASC11]. In particular we considered $n = 200$ eigenvalues of the Laplace-Beltrami operator, a variance $\sigma = 6(\phi_2 - \phi_1)$ and we evaluate at $M = 500$ values of energy e , where $e_{\min} = \log \phi_1 + 2\sigma$ and $e_{\max} = \log \phi_n - 2\sigma$ and ϕ_i denotes the i th eigenvalue of the Laplace-Beltrami operator.

Once local shape descriptors are computed, sparse coding is employed for local-to-global descriptor. For the numerical solution of the optimization problem (1) and (3) we use SPArse Modeling Software (SPAMS), an open-source optimization toolbox based on [MBPS09, MBPS10]. In all our experiments we make the trivial choice $\lambda = 1/2$, and learned

a dictionary D^c of $L = 500$ “words” from each class, such that D is a $M \times L = 500 \times 500$ matrix, where M is the length of each WKS.

Finally we compare the performances of our method with a global signature, namely Shape DNA [RWP06], and with the well known quantization approach [FKMS05], that we dubbed here as *Signature Distance Distribution* (SDD).

Shape DNA signature [RWP06] consists of the truncated spectrum of the Laplace-Beltrami operator. For the application of this popular global descriptor to retrieval scenarios, we follow the suggestions reported in [LGB*11]. More specifically, we consider only the first 13 eigenvalues and rescale the spectrum by the shape’s area to obtain the scale invariance of the descriptor.

Signature Distance Distribution (SDD) is a way to quantify a local descriptor for obtaining a global one. Once extracted local descriptors from the shape, the central idea of this method is to exploit the obtained informations to build up an histogram, that plays the role of a global descriptor. In particular, the histogram takes into account the occurrences of Euclidean distances of local signatures between every pair of random points on the shape. In order to capture the underlying geometry, the random selection of points are repeated several times. In this case, the random selection was repeated 10 times, in order to have about 10^4 distances between local descriptors. The output is a histogram ables to discriminate between different shapes, as reported in Figure 4.




Figure 4: Comparison between the SDD of the shapes of a cat and a dog. In the former there is a peak approximately around the 20th bin, in the latter around the 40th bin.

For matching purposes, at each pair of shapes of the same deformation category, we compute the ℓ^2 error between histogram’s occurrences vectors and consider correct the matching with the minimum error. Since the method involves a random procedure, Table 1 reports the mean of 5 runs of the algorithm.




Figure 2: Null shapes of our database taken from TOSCA [BBK08] and Sumner [SP04] databases. From left to right we find man, dog, cat, man, woman, horse, camel, cat, elephant, flamingo, horse and lioness. In all our experiment we consider correct the matching between two men, cats or horses. As you can see, the camel and elephant shapes are created by pose transfer from the galloping horse, and the lioness from pose transfer from the crouching cat.




Figure 3: Examples of deformations types considered in our database, taken from SHREC 11 robust benchmark. From left to right we find the null shape, affine, holes, micro holes, scale, sampling, noise, shot noise, topology and view.

Comparison results are shown in Table 1. The proposed method improves on all the deformation classes with respect to SDD or Shape DNA. In particular it improves drastically in the class of affine, big holes, topology and view deformations. It is worth noting that on view deformation the improvement with respect to Shape DNA was expected: in fact a global descriptor fails to identify correctly partial views of a shape. It is rather more interesting to observe that with this kind of deformation our method clearly outperform SDD. In noise and shot noise deformations our method performs like Shape DNA although sparse coding approach consider significantly more eigenvalues. It is a well known fact that the firsts eigenvalues is related to shape's lower-frequency contents, meanwhile higher eigenvalues is related to higher-frequency contents and manifest themselves as rough geometric features, i.e. shape details. Overall, the proposed method shows a clear improvement over other methods, by evidencing a more stable and robust behavior.

We notice also that our descriptor takes into account the behavior of an entire class of deformations and this should affect negatively the nearest neighbor performances. Indeed,

we believe that in nearest neighbor scenarios a shape descriptor behave better than an entire descriptor of an entire class of non-rigid deformations. The fact that results in Table 1 are very promising, give us an idea of the goodness of the method.

Figure 5 reports the dissimilarity matrix of the class of noise deformations. Note that this matrix is not computed as a classical dissimilarity matrix between shape descriptors. In fact our method generate a class descriptor. Each column represents a class of shape, in accordance to the shape representation shown in figure above the matrix. We have omitted the second instances of repeated classes as man, cat and horse for a better visual result. Each row represents the mean error between 5 non-rigid deformation and the underlying null shape with respect to the dictionary of the respective classes. In this particular case we have considered noise deformation. Blue colors represent small error values, red colors represent high error values. It is interesting to note that as expected man and woman classes has small error in comparison to other classes. This remark the *similarity* property of the proposed descriptor.

Table 1: Comparison between the Nearest Neighbor retrieved shape by SDD, Shape DNA and sparse coding approach.

Deformation	SDD		Shape DNA		sparse coding	
	corrects	%	corrects	%	corrects	%
Affine	45/72	0.67	49/72	0.68	59/72	0.82
Holes	36/72	0.50	58/72	0.81	65/72	0.90
Micro holes	65/72	0.90	64/72	0.89	65/72	0.90
Scale	72/72	1.00	71/72	0.99	72/72	1.00
Sampling	67/72	0.93	69/72	0.96	72/72	1.00
Noise	71/72	0.99	72/72	1.00	72/72	1.00
Shot noise	70/72	0.97	72/72	1.00	72/72	1.00
Topology	58/72	0.80	55/72	0.76	72/72	1.00
View	11/72	0.15	16/72	0.22	49/72	0.64
Average	495/648	0.76	526/648	0.81	598/648	0.92




Figure 5: Class-signature dissimilarity matrix for noise deformations. Blue colors represent lower values, red colors represent higher values.

Finally, in Figure 6 is reported an embedding in a 2D plane of the proposed local-to-global signatures. The embedding was performed through a multi-dimensional scaling algorithm (MDS). This figure highlights the good similarity properties of the proposed descriptor: man embedded signature is close to woman, and the same happens for cat and dog embedded signatures.




Figure 6: Embedding local-to-global signatures in a plane through MDS.

5. Conclusions

In this paper a new approach for local-to-global shape description is proposed. We have shown that sparse coding methods are particularly suitable to compactly describe a large set of point-based descriptors. Although we use in our experiment Wave Kernel Signature (WKS) only, our method is versatile in encoding any other local descriptors in order to inherit at the global level the desired properties of local behaviour.

We have evaluated our approach on 3D shape retrieval in the context of robustness against several shape deformations. Our approach has shown its effectiveness in dealing with such deformations by drastically improve state-of-the-art methods. In particular, thank to the sparsity constraint our method is able to detect the most relevant information of a given class of shapes and to ignore the irrelevant or confusing aspects which effect the correct shape retrieval.

Future work will be address on the evaluation of more advanced sparse coding methods in order to exploit discriminative learning in shape retrieval domain.

Acknowledgment

We would like to thank Alex and Michael Bronstein for useful suggestions and fruitful discussions.

References

- [ASC11] AUBRY M., SCHLICKWEI U., CREMENS D.: The Wave Kernel Signature: A quantum mechanical approach to shape analysis. In *Proc. of ICCV Workshop on Dynamic Shape Capture and Analysis (4DMOD)* (2011). 1, 2, 3, 5
- [BBB*11] BOYER E., BRONSTEIN A. M., BRONSTEIN M. M., BUSTOS B., DAROM T., HORAUD R.: SHREC 2011: robust feature detection and description benchmark. *Proc. of Eurographics Workshop on 3D Object Retrieval (3DOR)* (2011). 2
- [BBGO11] BRONSTEIN A. M., BRONSTEIN M. M., GUIBAS L. J., OVSJANIKOV M.: Shape Google: Geometric words and expressions for invariant shape retrieval. *ACM Transaction on Graphics (TOG)* 30, 1 (2011), 1–20. 1, 2
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical geometry of non-rigid shapes*. Springer, 2008. 5, 6
- [BBM*10] BRONSTEIN A. M., BRONSTEIN M. M., MAHMOUDI M., KIMMEL R., SAPIRO G.: A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision* 89, 1 (2010), 56–68. 2
- [BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signature for non-rigid shape recognition. In *Proc. of Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 1704–1711. 1, 2, 3
- [BMP02] BELONGIE S., MALIK J., PUZICHA J.: Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 4 (2002). 2
- [Bro] BRONSTEIN A. M.: Spectral descriptors for deformable shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)* (submitted). 2, 3, 4
- [CCM11] CASTELLANI U., CRISTANI M., MURINO V.: Statistical 3D shape analysis by local generative descriptors. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)* 33 (2011), 2555–2560. 1, 2
- [DK12] DAROM T., KELLER Y.: Scale-invariant features for 3D mesh models. *IEEE Trans. on Image Processing* 21 (2012), 2758–2769. 1, 2
- [FHK*04] FROME A., HUBER D., KOLLURI R., BULOW T., MALIK J.: Recognizing objects in range data using regional point descriptors. *Lecture Notes in Computer Science* (2004), 224–237. 2
- [FKMS05] FUNKHOUSER T., KAZHDAN M., MIN P., SHILANE P.: Shape-based retrieval and analysis of 3D models. *Communications of the ACM* 48 (2005), 58–64. 1, 2, 5
- [FMK*03] FUNKHOUSER T., MIN P., KAZHDAN M., CHEN J., HALDERMAN A., DOBKIN D., JACOBS D.: A search engine for 3D models. *ACM Transactions on Graphics (TOG)* 22 (2003), 83–105. 1
- [IJL*05] IYER N., JAYANTI S., LOU K., KALYNARAMAN Y., RAMANI K.: Three-dimensional shape searching: state-of-the-art review and future trends. *Computer-Aided Design* 37 (2005), 509–530. 1
- [JH99] JOHNSON A., HEBERT M.: Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21, 5 (1999), 433–449. 2
- [KCB09] KHOVALED S., CASTELLANI U., BARTOLI A.: Semantic shape context for the registration of multiple partial 3--D views. In *BMVC* (2009). 2
- [LGB*11] LIAN Z., GODIL A., BUSTOS B., DAOUDI M., ET AL.: SHREC 2011 track: shape retrieval on non-rigid 3D watertight meshes. In *Proc. of Eurographics Symposium on 3D Object Retrieval* (2011). 1, 2, 5
- [LGF*10] LIAN Z., GODIL A., FABRY T., T. F., ET AL.: SHREC 2010: Non-rigid 3D shape retrieval. In *Proc. of Eurographics Workshop on 3D Object Retrieval (3DOR)* (2010). 1
- [MBPS09] MAIRAL J., BACH F., PONCE J., SAPIRO G.: Online dictionary learning for sparse coding. In *Proc. of the International Conference on Machine Learning (ICML)* (2009), pp. 689–696. 1, 2, 3, 4, 5
- [MBPS10] MAIRAL J., BACH F., PONCE J., SAPIRO G.: Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research* 11 (2010), 19–60. 1, 2, 3, 4, 5
- [MGGP06] MITRA N. J., GUIBAS L., GIESEN J., PAULY M.: Probabilistic fingerprints for shapes. In *Symposium on Geometry Processing* (2006), pp. 121–130. 2
- [PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Computer Graphics Forum (Eurographics) (to appear)* (2013). 2, 3
- [RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-Beltrami spectra as ‘Shape-DNA’ of surfaces and solids. *Computer-Aided Design* 38 (2006), 342–366. 2, 5
- [SF06] SHILANE P., FUNKHOUSER T.: Selecting distinctive 3D shape descriptors for similarity retrieval. In *Proc. of IEEE International Conference on Shape Modelling and Applications* (2006). 1
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. of the Symposium on Geometry Processing* (2009), pp. 1383–1392. 1, 2, 3
- [SP04] SUMNER R. W., POPOVIĆ J.: Deformation transfer for triangle meshes. *ACM Trans. on Graphics (TOG)* 23 (2004), 399–405. 5, 6
- [SS02] SCHÖLKOPF B., SMOLA A. J.: *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, 2002. 2
- [TCF10] TOLDO R., CASTELLANI U., FUSIELLO A.: The bag of words approach for retrieval and categorization of 3D objects. *The Visual Computer: International Journal of Computer Graphics* 26 (2010), 1257–1268. 1, 2
- [Tib96] TIBSHIRANI R.: Regression shrinkage and selection via the Lasso. *J. Royal Stat. Society* (1996), 267–288. 1, 3
- [TV04] TANGELDER J. W., VELTKAMP R. C.: A survey of content based 3D shape retrieval methods. In *International Conference on Shape Modelling and Applications* (2004), pp. 145–156. 1, 2