

Research Article

Investigating the Bag-of-Words Method for 3D Shape Retrieval

Xiaolan Li¹ and Afzal Godil²

¹ College of Computer Science & Information Engineering, Zhejiang Gongshang University, Hangzhou, Zhejiang 310018, China

² Information Technology Laboratory, National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

Correspondence should be addressed to Xiaolan Li, lixiaolana@hotmail.com

Received 1 December 2009; Revised 27 February 2010; Accepted 3 March 2010

Academic Editor: Yingzi Du

Copyright © 2010 X. Li and A. Godil. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

This paper investigates the capabilities of the Bag-of-Words (BWs) method in the 3D shape retrieval field. The contributions of this paper are (1) the 3D shape retrieval task is categorized from different points of view: specific versus generic, partial-to-global retrieval (PGR) versus global-to-global retrieval (GGR), and articulated versus nonarticulated (2) the spatial information, represented as concentric spheres, is integrated into the framework to improve the discriminative ability (3) the analysis of the experimental results on Purdue Engineering Benchmark (PEB) reveals that some properties of the BW approach make it perform better on the PGR task than the GGR task (4) the BW approach is evaluated on nonarticulated database PEB and articulated database McGill Shape Benchmark (MSB) and compared to other methods.

1. Introduction

With recent advances in scanning and modeling technologies, large number of 3D models are created and stored in databases. For these databases be used effectively require methods for indexing, retrieval, and clustering. Therefore, retrieval and classification of 3D objects are becoming an increasingly important task in modern applications such as computer vision, computer aided design/computer aided manufacturing, multimedia, molecular biology, biometric, security, and robotics.

Because of its simplicity, flexibility, and effectiveness, the Bag-of-Words (BWs) method, which originated from the document retrieval field, has recently attracted large amount of interests in the computer vision fields. It has been applied in the applications such as image/video classification [1], 3D shape analysis, and retrieval [2–5]. We will explore its performance especially for the 3D shape retrieval task in this paper.

A typical 3D shape retrieval task can be defined as: giving a query 3D shape, to obtain a list of 3D shapes ordered by the similarity between the query object and the one on the list. Several methods are proposed to solve the problem, such as Light Field descriptors [6], spherical harmonics descriptor [7], D2 shape distribution [8], Reeb Graph-based descriptors [9], Local Feature-based methods [4, 5]. The performance of the methods varies mainly according to the specific tasks. In

fact, from different points of view, the 3D shape retrieval task can be further refined as follows.

(1) Differentiated from the object category extent, the task can be discussed in “specific” and “generic” domain, which depends on the purpose and interest of the specialists. The representative benchmarks of the latter one include Princeton Shape Benchmark [10], NIST 3D Benchmark [11], while CAD [12], Protein [13], and Biometrics [14] analysis are several important “specific” domains, which have their own properties. For example, CAD models have more complicated structure with holes and other local features. Only using global information, these subtle details could be neglected and lead to less ideal retrieval results.

(2) Based on the completeness of the query shape, the task can be divided into two subtasks as “Partial-to-Global Retrieval (PGR)” and “Global-to-Global Retrieval (GGR)”. For the former one, every query shape is regarded as an incomplete object, which is used to obtain similar complete objects from the database. This happens in many cases. For example, when using the 3D range scanners to capture 3D data in real time, because of the limitation of the view angle, the occlusion in the scene, and the real time requirement, only parts of the object can be captured during scanning. Then this incomplete point clouds may be used as the query shape to retrieve the corresponding complete model from an existing database. Solving this problem will also benefit

several other applications, such as data registration [15] and model fixing [16]. Most of the global-based shape retrieval methods [6, 8], which require the complete geometry of a 3D object, cannot be applied directly to PGR. To our knowledge there are only a few literature [3, 17] contributions that solve the PGR problem.

(3) Based on the deformability of the shape, there exist “Articulated Shape Retrieval (ASR)” and “Nonarticulated Shape Retrieval (NASR)”. Lots of the natural and man-made objects are deformable. For instance, in CAESAR [14], each person is scanned in three different postures: standing, sitting with arms open, and sitting with arms down. When performing shape retrieval using a sitting model of person A as the query model, the preferred result is to obtain the other two different gestured model of person A than to retrieve the sitting models of other persons. According to the results in [4], Light Field method [6], which performs greatly when dealing with NASR problem, produces poor results for ASR task.

To some extent, in [3, 4, 17], the above three different tasks are discussed within the BWs framework, but there still lacks thorough investigation. Several open problems remain unsolved, such as how to integrate spatial information into the BWs framework to improve the performance. In this paper, we investigate deeply into these three different cases within the framework of the BWs method with spin images [18] as the low-level features, and provide profound experimental results to support the discussion.

The organization of the paper is as follows. Several related works are summarized in Section 2. The performance measure is discussed in Section 3. In Section 4, we first introduce the ordinary procedure for BWs framework in 3D domain. Then take the CAD database [12] as an example, a Concentric Bag-of-Words (CBWs) approach is proposed to enhance the discriminative ability of the original BWs method. Several interesting phenomena are studied for PGR problem in Section 5. As for ASR task, McGill articulated shape benchmark [19] is adopted to test the effectiveness of our approach in Section 6. Finally, we conclude the paper in Section 7.

2. Related Work

Many efforts have been taken to perform 3D shape retrieval recently. Among them, the BWs method, which represents a 3D shape as an orderless collection of local features, has demonstrated impressive level of performance.

In [2, 3], BWs method is explored to accomplish PGR task, in which a visual feature dictionary is constituted by clustering spin images [18]. Then, Kullback-Leibler divergence is proposed as a similarity measurement in [3], while a probabilistic framework is introduced in [2].

For the ASR task, Ohbuchi et al. [4] apply the SIFT algorithm to depth buffer images of the model captured from uniformly sampled locations on a view sphere to collect visual words. After vector quantization, Kullback-Leibler divergence measures the similarities of the models. It also demonstrates that a) given enough samples, the BWs method can reach a comparable retrieval result as a vision

based method like Light Field [6], when dealing with NASR task; b) the BWs method performs better than Light Field [6] when dealing with ASR task. In this paper, spin images are used as local features, which can be directly extracted as many as you want in 3D domain. On the other hand, according to [1], dense features, such as spin images, perform better than sparse features, such as SIFT.

Although the BWs method has many advantages, it suffers from its lack of spatial information. Some methods focus on integrating the spatial layout information into the BWs method. Lazebnik et al. [20] proposes a spatially enriched Bags-of-Words approach. It works by partitioning the image into increasingly fine subregions and computing histograms of local features found inside each subregion. Implicitly geometric correspondences of the subregions are built in the pyramid matching scheme [21]. In [22], the object is an ensemble of canonical parts linked together by an explicit homographic relationship. Through an optimization procedure, the model, corresponding to the lowest residual error, gives the class label to the query object along with the localization and pose estimation. Yuan and Wu [23] describes a context aware clustering method, which captures the contextual information between data. For the BWs method, it means the visual dictionary is constructed based on both the primitive visual features and spatial contexts. Li et al. [5] propose to treat the model in two different domains, named the feature domain and the spatial domain. The visual word dictionary is built in the feature domain as in the ordinary BWs method. On the other side, the whole model is partitioned into several pieces in the spatial domain. Thereafter, each piece of the model is represented as a word histogram. The whole model is recorded as several word histograms along with a geometry matrix which stores the relative distances between every pairs of the pieces. The weighted sums of dissimilarity measurements from these two domains are used to measure the differences between models.

3. Performance Measure

The performance measure used in this study is the precision-recall curve. Precision-recall curve is the most common metric to evaluate 3D shape retrieval system. Precision is the ratio of retrieved objects that are relevant to all retrieved objects in the ranked list. Recall is the ratio of relevant objects retrieved in the ranked list to all relevant objects.

Let A be the set of all relevant objects, and B be the set of all retrieved objects, then

$$\text{precision} = \frac{A \cap B}{B}, \quad \text{recall} = \frac{A \cap B}{A}. \quad (1)$$

Basically, recall evaluates how well a retrieval algorithm finds what we want and precision evaluates how well it weeds out what we do not want. There is a tradeoff between recall and precision, one can increase recall by retrieving more, while can decrease precision.

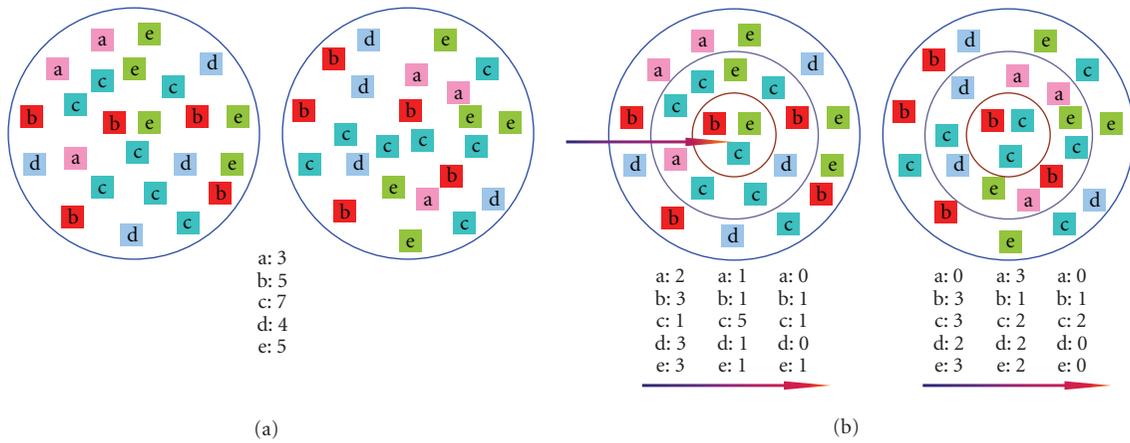


FIGURE 1: Comparing BWs and CBW representations. (a) Representing shapes with one global *Bag-of-Words* model. The left and the right shapes are both composed with 5 different words: *a, b, c, d, e*. Both feature vectors are $[3, 5, 7, 4, 5]$, which count the occurrences of each word. That means using BWs representation, the left and the right shapes are regarded as the same. (b) Representing shapes with *Concentric Bags-of-Words* model. Even there are the same two shapes as shown in (a), because of the concentric sphere partitioning, the left and the right shapes are different. Along the arrow's direction, counting from the outer sphere to the inner one, their feature vectors are $[2\ 3\ 1\ 3\ 3; 1\ 1\ 5\ 1\ 1; 0\ 3\ 3\ 2\ 3; 3\ 1\ 2\ 2\ 2; 0\ 1\ 2\ 0\ 0]$, respectively.

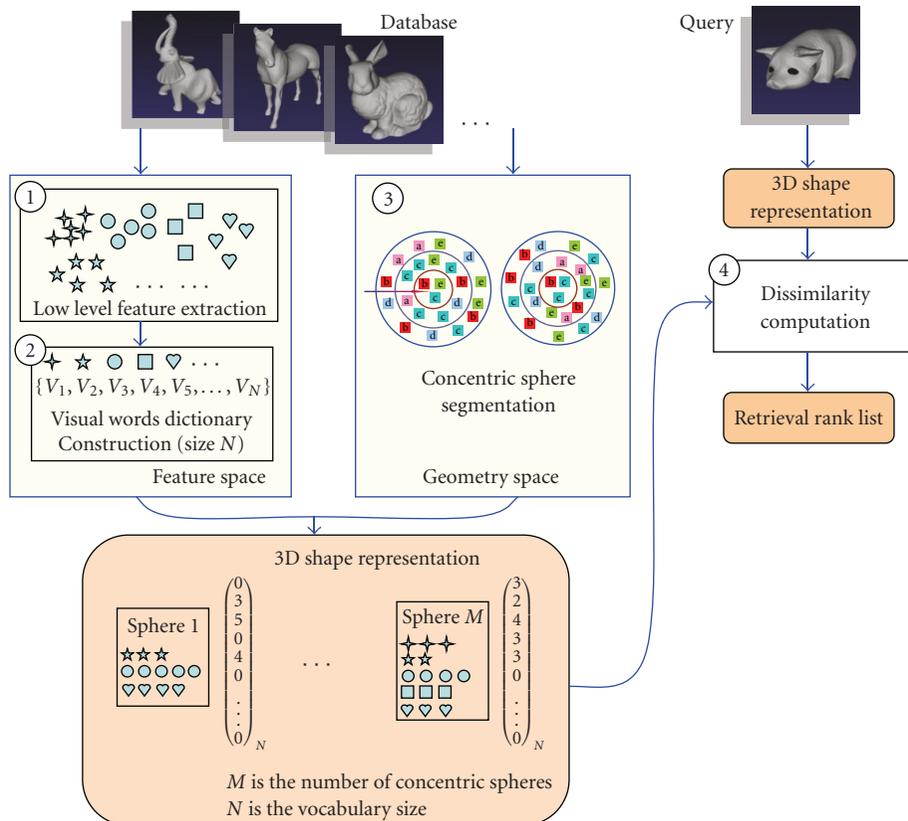


FIGURE 2: A schematic description of CBW method.

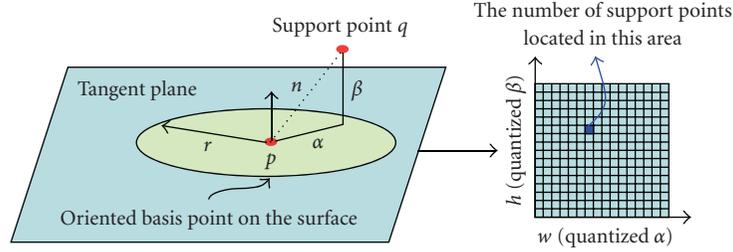


FIGURE 3: The demonstration of spin image.

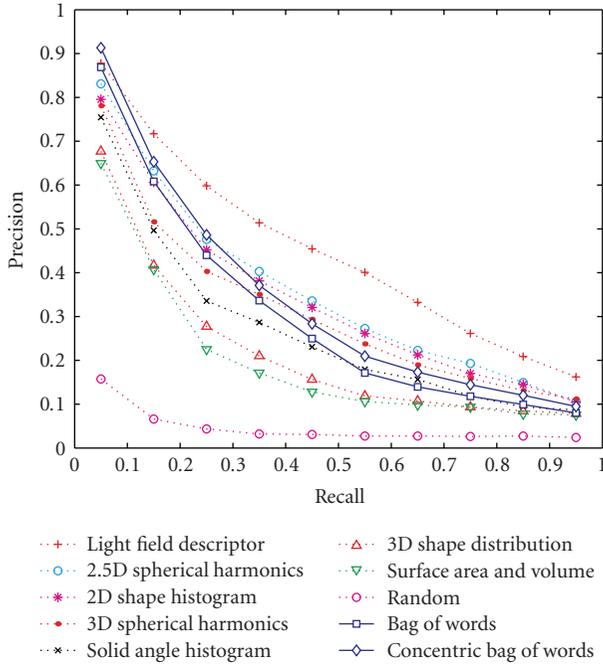
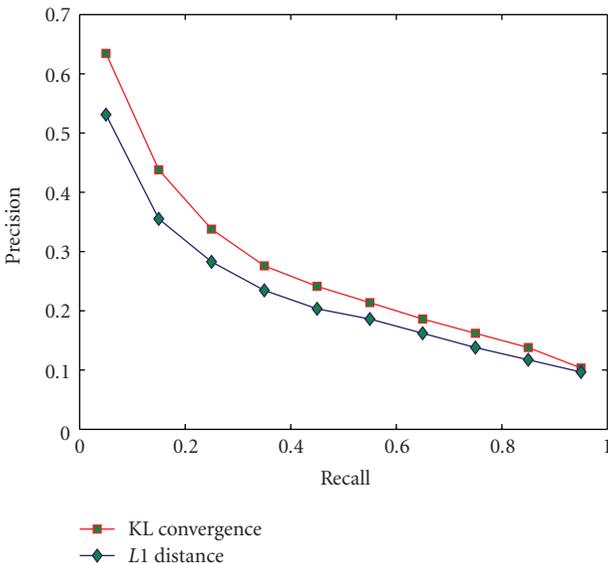


FIGURE 4: Precision-recall (PR) plots of CBW, BWs, and other methods listed in [12].

FIGURE 5: Precision-recall (PR) plots obtained with KL and L_1 dissimilarity measurement when doing PGR.

4. Bag-of-Words and Concentric Bag-of-Words Methods

We first describe the original formulation of BWs representation [1, 3], and then introduce the whole procedure of Concentric Bags-of-Words (CBW) method. Their difference is demonstrated in Figure 1. Its effectiveness is shown by the experiment performed on “specific” shape database PEB [12] to reveal its effectiveness.

4.1. Bag-of-Words Descriptor. Let us use the ordinary 3D shape retrieval as an example to give an explanation of the BWs framework. Denote N be the total number of labels (“visual words”) in the learned visual dictionary. The 3D shape can be represented as a vector with length N , in which the elements count the occurrences of the corresponding label. The procedure can be completed in three steps.

- (1) Local feature descriptors, such as spin image [18], are applied to the 3D model to acquire low-level features.
- (2) *Visual words*, denoted as the discrete set $\{V_1, V_2, \dots, V_N\}$, are formed by clustering the features into N clusters, so that each local feature is assigned to a discrete label.
- (3) The shape of the 3D model is summarized with a global histogram (“*Bag-of-Words*”), denoted as a vector $f_v = (x_1, x_2, \dots, x_N)$, by counting the occurrences of each *visual word*.

4.2. Concentric Bag-of-Words Method. Rather than using only a global histogram, this paper advocates using more than one histogram along with its related spatial information to reveal the 3D shape in more detail. Specifically, the model is partitioned with several concentric spheres, and all the parts between two neighboring spheres are recorded with original BWs descriptor, which leads to the name Concentric Bag-of-Words. A schematic description of the approach is given in Figure 2.

The first block in Figure 2 represents low-level feature extraction. Although several local features, such as depth buffer image, can be adopted to extract low-level features from 3D models, spin image is the one adopted here. Compared to using depth buffer image, adopting spin image has at least two advantages. First, it is quicker to compute. Second, it can capture the details from the concave

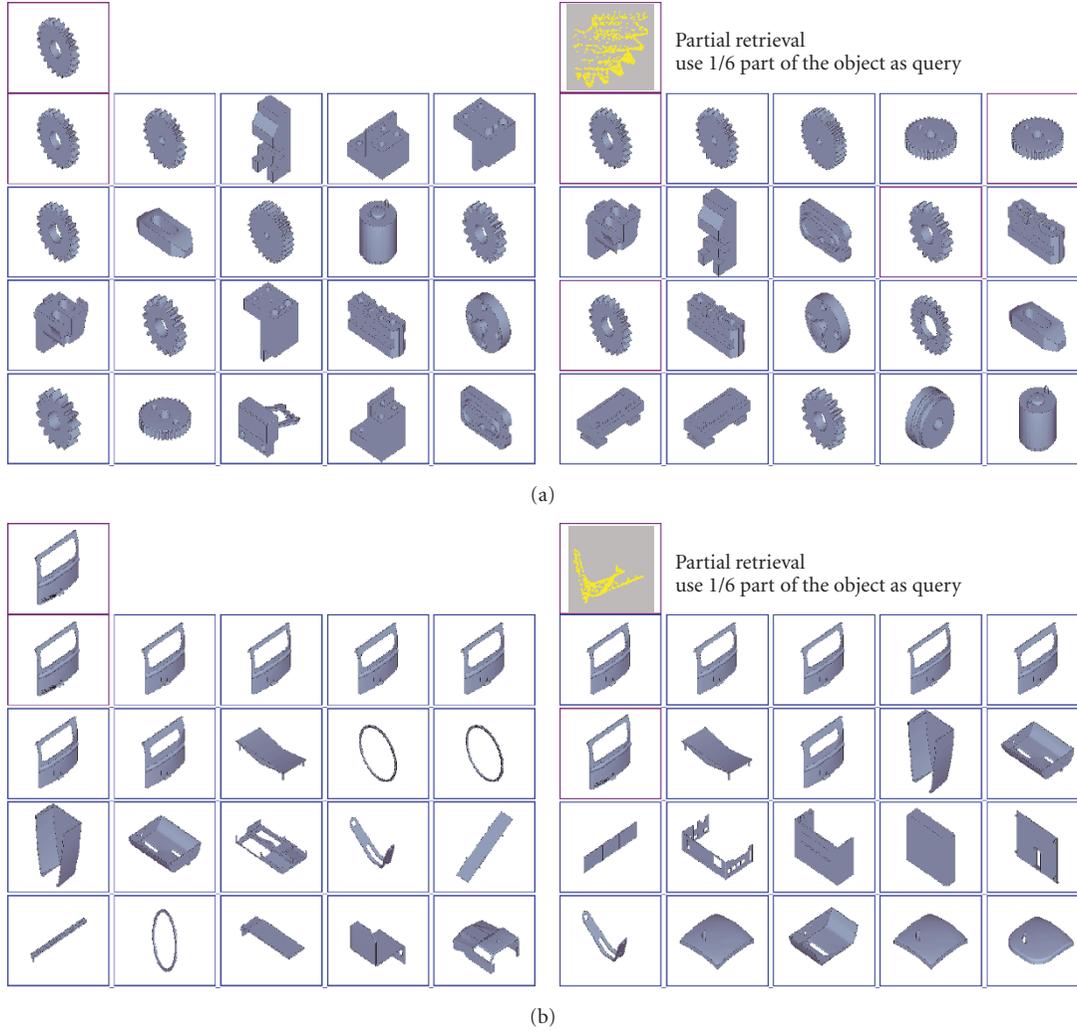


FIGURE 6: The example to show the difference between GGR and PGR. (a) First example to show the difference between Global-to-Global retrieval (GGR) and Partial-to-Global retrieval (PGR). The left group shows the GGR result using a complete model (the top-left image) as the query. The right group shows the PGR result using 1/6 part of the complete model (the small image shown in the bottom-left corner of the first image) as the query. The top 20 models are listed orderly according to the dissimilarity measurement. (b) The second example to show the difference between GGR and PGR. The layout of the images is the same as that of (a).

area and the self-hidden area. As shown in Figure 3, it characterizes the local properties around its basis point p within the support range r . It is a two-dimensional histogram accumulating the number of points located at the coordinate (α, β) , where α and β are the lengths of the two orthogonal edges of the triangle formed by the oriented-basis point p , whose orientation is defined by the normal n , and support point q . The final size of the spin images is defined by the width w and the height h of the spin plane. We uniformly sample N_b oriented-basis points and N_s support points on the surface of the model, which satisfies insensitivity to the tessellation and resolution of the mesh.

After extracting a set of spin images for each model, we construct a shape dictionary as shown in the second block, whose size is predetermined as N , by clustering all spin images acquired from the whole training dataset with k -means method.

Instead of representing one model with a histogram of the words from the dictionary, it is partitioned into M regions by grouping the oriented-basis points with M concentric spheres as demonstrated in the third block. Thereafter, the model is recorded as a set of histograms.

Because all the models are scaled into unified scale and the partitioning is also unified, the correspondence between the regions of two models is obvious. It can be constructed from outer sphere to inner sphere, as shown in Figure 1(b), or reverse. Thus, the CBW feature vector is recorded as

$$\begin{aligned}
 cfv &= (fv^1, fv^2, \dots, fv^M) \\
 &= (x_1^1, x_2^1, \dots, x_N^1, \dots, x_1^M, x_2^M, \dots, x_N^M).
 \end{aligned} \tag{2}$$

When performing 3D shape retrieval, the CBW representation of the query shape is constructed on line, and

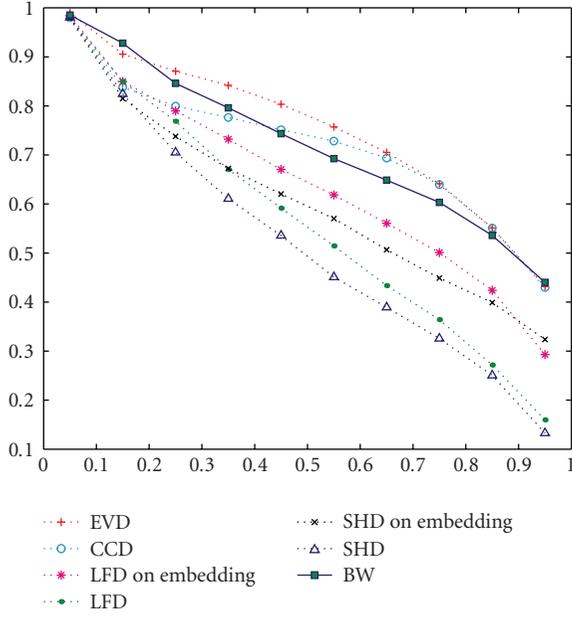


FIGURE 7: Precision-recall (PR) plots for various descriptors when applied to the McGill database of articulated shapes [19]. Except BWs, all of the other results can be found in [25].

compared with those in the database. An ordered retrieval list is obtained according to the dissimilarity metric, which is

$$\text{Dist}(O^A, O^B) = \sum_{i=1}^M \text{dist}(cfv_i^A, cfv_i^B), \quad (3)$$

where O^A , O^B are two objects A and B , respectively; $\text{dist}(cfv_i^A, cfv_i^B)$ measures the dissimilarity between two feature vectors, which can be KL divergence [3], cosine distance, $L1$ and $L2$ distance. Thus, for every query object, the objects in the database are all assigned a metric value based on (3), which results a sorted retrieval list.

4.3. Experimental Results. According to the discussions in [3, 17], several parameters related to the CBW approach are defined as follows.

- (1) The support range r : $r = 0.4 * R$, where R is the radius of the model.
- (2) The width w and the height h of the spin plane: $w = h = 12$.
- (3) The number of oriented-basis points for one model N_b : $N_b = 500$.
- (4) The number of oriented-basis points for one model N_s : $N_s = 5000$.
- (5) The size of the dictionary N : $N = 1500$.
- (6) The number of the concentric spheres M : $M < 10$.

The CBW approach can be applied both in “specific” and “generic” domains. Here we demonstrate it on Purdue

Engineering Benchmark [12], which contains 866 3D CAD models and is classified into 42 classes such as, “Discs”, “T-shaped parts”, and “Bracket-like parts”. In Figure 4 we compare the Precision-Recall curves obtained with CBW and BWs using $L1$ as dissimilarity measurement to those methods defined in [12], such as Light Field Descriptor, 2.5D Spherical harmonics, 2D Shape Histogram, 3D Spherical Harmonics, Solid Angle Histogram, 3D Shape Distribution, Surface Area and Volume. Here, for CBW, $M = 9$. Obviously, the concentric sphere partition improves the PR rate, and makes the local Feature-based method comparable to the global Feature-based method, such as 2.5D Spherical Harmonics listed as the second best method in [12].

5. Partial-to-Global Retrieval

In CAD domain, the PGR task is specifically important. Suppose that the query partial model is a screw, the target complete model we want to obtain is the same size screw with screw cap on. Here, we design an experiment on PEB to simulate the case described above: (1) Represent the models with BWs method and save the descriptors for the following usage as block 1 and 2 in Figure 2. (2) When performing the Partial-to-Global Retrieval, the sampled oriented-basis points are grouped into M_p regions according to their geometric positions first. Then one of the groups is chosen as the partial query shape. The BWs representation is constructed on line and used to compare with the saved BWs descriptors of the complete models. The requirements for dissimilarity measure for the partial-to-global retrieval task are quite different from the global-to-global retrieval problem. As described in [3], the dissimilarity between the query data and the target model is not equal to that between the target model and the query data. It means that the dissimilarity metric should be asymmetry. An ordinary symmetric distance measurement, such as $L1$, $L2$, is not a suitable choice. KL divergence is chosen here to satisfy the asymmetric property. When using one sixth of the model to be the query shape, the two PR curves in Figure 5 demonstrate the improvement introduced by KL comparing to $L1$.

Figure 6 provides two examples comparing the retrieval results of Global-to-Global Retrieval and Partial-to-Global Retrieval, in which one sixth of a gear is used as the query shape. It shows that the PGR is better than the GGR, since PGR lists more gears on the top of the list than GGR does. Why does PGR perform better? Recalling the definition of the feature vector will provide some clues to the answer. The feature vector describes the frequency of the visual words appearing in the shape. When using the entire gear model to be the query data, the plane-kind of visual word overwhelm the other features. However, using partial of the object to be the query data, the gear teeth shape dominates the whole shape. So more gears are picked out and listed on the top of the list.

6. Articulated Shape Retrieval

The Articulated Shape Retrieval requires that the shape descriptor should be deformation invariant, which is not

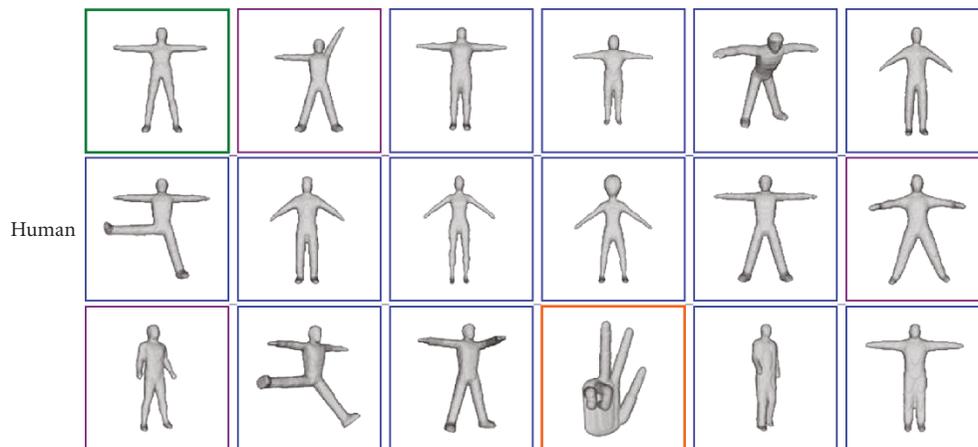
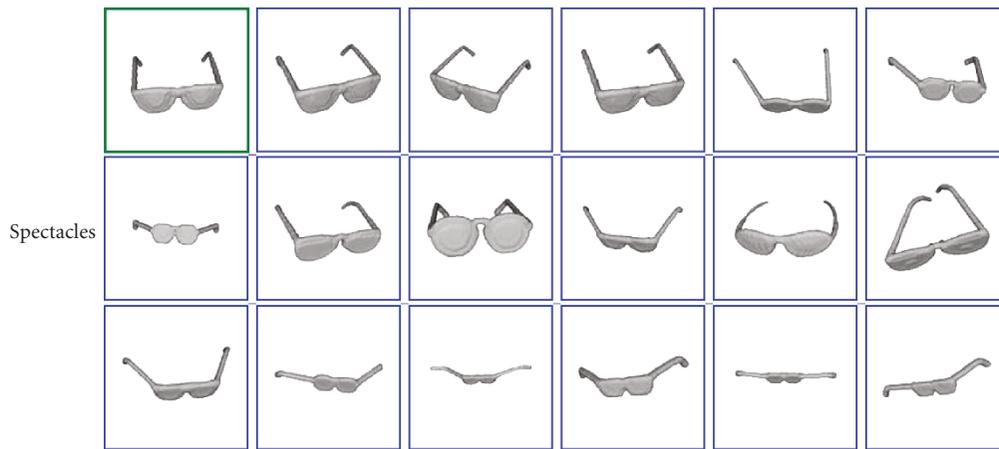
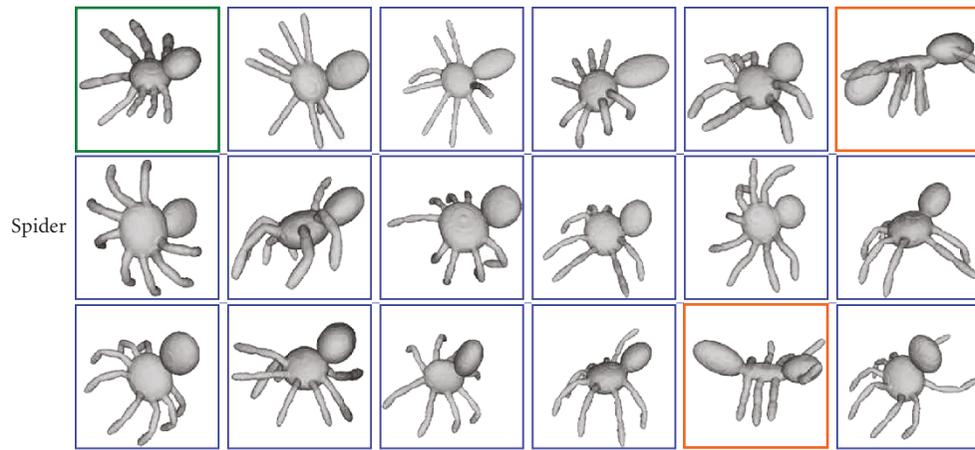


FIGURE 8: Some retrieval results from the McGill database [19]: green bold frame defines the query shape, and orange bold frame defines the false pick up. Three groups show three different shapes, in which from top to bottom they are spider, spectacles, and human.

satisfied by several previous methods [6–8]. They perform well when dealing with rigid objects, but manifest poor performance when dealing with deformable ones [4]. BWs method can still be used effectively for ASR task. The descriptors for the models are constructed by following the procedure shown as block 1 and 2 in Figure 2.

We applied the BWs method for ASR task on McGill Shape Benchmark (MSB) [19]. The configuration of the parameters is almost the same as those listed in Section 4.3, except that (a) the width w and the height h of the spin plane: $w = h = 16$, (b) the number of oriented-basis points for one model N_b : $N_b = 1000$. Since all of the models in MSB are regarded as complete, $L1$ distance is chosen to measure the dissimilarity.

In Figure 7, the BW-based retrieval result is compared with several methods described in [25]. BWs method is comparable to the best method EVD. However, except BWs, all the other methods are based on geodesic distance computation, which is computational expensive. On the contrary, our method is constrained on local area and can be applied for on-line retrieval.

Figure 8 shows three visual results of articulated shape retrieval. Only the top 18 results are listed here, in which the green bold framed shape is the query shape, and the orange bold framed shapes are the false recall. Figure 8(a) shows the results of retrieving a spider shape from the database. Among these 18 retrieval shapes, only two shapes do not belong to the spider class but the ant class. Figures 8(b) and 8(c) are the results using a spectacles and a human shape as query model, respectively. Even though there are quite large amount of bending in the shapes, the performance is quite good.

7. Conclusion and Discussion

In this paper, we explore the BWs framework to solve several different tasks in 3D shape retrieval field, which are classified as specific versus generic, partial-to-global versus global-to-global retrieval, and articulated versus Nonarticulated. For each type, the effectiveness of BWs method is discussed in detail. First, CBW method is introduced to improve the discrimination ability of original BWs representation. Second, BWs is applied on PEB to perform partial-to-global retrieval task. And several results revealed, for some shape (gear-like shape), that PGR performs better than GGR. Finally, we compared the results of BWs to several other methods on McGill articulated shape database. Our results are comparable to the best results in [25]. More experiments need to be done to verify the influence of the parameters listed in Section 4.3.

Acknowledgments

The authors would like to thank the SIMA program and the IDUS program for supporting this work. This work has also been partially supported by NSF Grants of China (60873218) and NSF Grants of Zhejiang (Z1080232).

References

- [1] L. Fei-Fei and P. Perona, "A bayesian hierarchical model for learning natural scene categories," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '05)*, vol. 2, pp. 524–531, June 2005.
- [2] Y. Shan, H. S. Sawhney, B. Matei, and R. Kumar, "Shapeme histogram projection and matching for partial object recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 4, pp. 568–577, 2006.
- [3] Y. Liu, H. Zha, and H. Qin, "Shape topics: a compact representation and new algorithms for 3D partial shape retrieval," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '06)*, vol. 2, pp. 2025–2032, 2006.
- [4] R. Ohbuchi, K. Osada, T. Furuya, and T. Banno, "Salient local visual features for shape-based 3D model retrieval," in *Proceedings of the IEEE International Conference on Shape Modeling and Applications (SMI '08)*, pp. 93–102, Stony Brook, NY, USA, 2008.
- [5] X. Li, A. Godil, and A. Wagan, "Spatially enhanced bags of words for 3D shape retrieval," in *Proceedings of the 4th International Symposium on Advances in Visual Computing (ISVC '08)*, vol. 5358 of *Lecture Notes in Computer Science*, pp. 349–358, Las Vegas, Nev, USA, 2008.
- [6] D.-Y. Chen, M. Ouhyoung, X.-P. Tian, and Y.-T. Shen, "On visual similarity based 3D model retrieval," *Computer Graphics Forum*, vol. 22, no. 3, pp. 223–232, 2003.
- [7] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, "Rotation invariant spherical harmonic representation of 3D shape descriptors," in *Proceedings of the ACM International Conference Symposium on Geometry Processing*, vol. 43, pp. 156–164, Aachen, Germany, June 2003.
- [8] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Shape distributions," *ACM Transactions on Graphics*, vol. 21, no. 4, pp. 807–832, 2002.
- [9] S. Biasotti, "Reeb graph representation of surfaces with boundary," in *Proceedings of the Shape Modeling International (SMI '04)*, pp. 371–374, 2004.
- [10] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser, "The princeton shape benchmark," in *Proceedings of the International Conference on Shape Modeling and Applications (SMI'04)*, pp. 167–178, Genova, Italy, June 2004.
- [11] R. Fang, A. Godil, X. Li, and A. Wagan, "A new shape benchmark for 3D object retrieval," in *Proceedings of the 4th International Symposium on Visual Computing*, vol. 5358 of *Lecture Notes in Computer Science*, pp. 381–392, Las Vegas, Nev, USA, 2008.
- [12] S. Jayanti, Y. Kalyanaraman, N. Iyer, and K. Ramani, "Developing an engineering shape benchmark for CAD models," *Computer Aided Design*, vol. 38, no. 9, pp. 939–953, 2006, Shape Similarity Detection and Search for CAD/CAE Applications.
- [13] H. M. Berman, J. Westbrook, Z. Feng, et al., "The protein data bank," *Nucleic Acids Research*, vol. 28, no. 1, pp. 235–242, 2000.
- [14] CAESAR Anthropometric Database, <http://store.sae.org/caesar/>.
- [15] N. J. Mitra, L. Guibas, J. Giesen, and M. Pauly, "Probabilistic fingerprints for shapes," in *Proceedings of the 4th Eurographics Symposium on Geometry Processing*, vol. 256, pp. 121–130, Sardinia, Italy, 2006.
- [16] T. Funkhouser, M. Kazhdan, P. Shilane, et al., "Modeling by example," in *Proceedings of the 31st International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '04)*, pp. 652–663, 2004.

- [17] X. Li, A. Godil, and A. Wagan, "3D part identification based on local shape descriptors," in *Proceedings of the Performance Metrics for Intelligent Systems Workshop (PerMIS '08)*, Gaithersburg, Md, USA, August 2008.
- [18] A. E. Johnson and M. Hebert, "Using spin images for efficient object recognition in cluttered 3D scenes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 5, pp. 433–449, 1999.
- [19] J. Winn, A. Criminisi, and T. Minka, "Object categorization by learned universal visual dictionary," in *Proceedings of the 10th IEEE International Conference on Computer Vision (ICCV '05)*, vol. 2, pp. 1800–1807, 2005.
- [20] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: spatial pyramid matching for recognizing natural scene categories," in *Proceedings of the Computer Vision and Pattern Recognition (CVPR '06)*, vol. 2, pp. 2169–2178, 2006.
- [21] K. Grauman and T. Darrell, "The pyramid match kernel: discriminative classification with sets of image features," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV '05)*, vol. 2, pp. 1458–1465, October 2005.
- [22] S. Savarese and L. Fei-Fei, "3D generic object categorization, localization and pose estimation," in *Proceedings of the IEEE 11th International Conference on Computer Vision (ICCV '07)*, pp. 1–8, Rio de Janeiro, Brazil, October 2007.
- [23] J. Yuan and Y. Wu, "Context aware clustering," in *Proceedings of the Computer Vision and Pattern Recognition (CVPR '08)*, pp. 1–8, Anchorage, Alaska, USA, June 2008.
- [24] R. Gal and D. Cohen-Or, "Salient geometric features for partial shape matching and similarity," *ACM Transactions on Graphics*, vol. 25, no. 1, pp. 130–150, 2006.
- [25] V. Jain and H. Zhang, "A spectral approach to shape-based retrieval of articulated 3D models," *Computer Aided Design*, vol. 39, no. 5, pp. 398–407, 2007.