

A Salient-Point Signature for 3D Object Retrieval

Indriyati Atmosukarto, Linda G. Shapiro
Department of Computer Science and Engineering
University of Washington
Seattle, WA, USA
indria, shapiro@cs.washington.edu

ABSTRACT

In this paper we describe a new 3D object signature and evaluate its performance for 3D object retrieval. The signature is based on a learning approach that finds the characteristics of salient points on a 3D object and represents the points in a 2D spatial map based on a longitude-latitude transformation. Experimental results show that the signature is able to achieve good retrieval scores for both pose-normalized and randomly-rotated object queries.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval; H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing

General Terms

Algorithms, Performance

Keywords

3D object retrieval, 3D object signature, salient points

1. INTRODUCTION

With the recent advancement in technology for digital acquisition of 3D models, there has been an increase in the availability and usage of 3D objects in a number of fields, including medical applications, architectural research, and engineering research. As a result, there is a large collection of 3D objects available. This motivates the need to be able to retrieve 3D objects that are similar in shape to a given 3D object query. Current techniques for text retrieval, 2D image retrieval, and video retrieval cannot be directly translated and applied to 3D object shape retrieval, as 3D objects have different data characteristics from other data modalities.

Shape-based retrieval of 3D objects is an important research topic. The accuracy of a shape-based retrieval system largely depends on finding a good descriptor that is able to

represent the local and global characteristics of the 3D object. In this paper, we describe our method of representing a 3D object and its use in a content-based retrieval system. We use a learning approach to identify the interesting local features or salient points on a 3D object, and we represent these feature points in a 2D spatial map. Retrieval of 3D objects is then performed by comparing the 2D spatial map of a query to the 2D maps of all the objects in the database and retrieving the database objects in order of the similarity of their 2D maps to the query.

The rest of this paper is organized as follows. We first discuss existing shape descriptors and their limitations. We then describe our data acquisition process. Next we describe our method for learning the salient points of a 3D object. We then describe a 2D longitude-latitude map signature that captures the pattern of the salient points. Finally, we explain and analyze the results from our experiments and provide a summary and suggestions for future work.

2. RELATED WORK

Content-based 3D object retrieval has received increased attention in the past few years due to the increased number of 3D objects available. There have been several survey papers on the topic [15, 4, 6], and recently researchers in the area have taken the initiative to organize an annual 3D shape retrieval evaluation called the SHREC - 3D Shape Retrieval Contest, whose general objective is to evaluate the effectiveness of 3D shape retrieval algorithms. Del Bimbo et al. [3] presented a comparative analysis of a number of different descriptors in retrieving 3D models from digital archives. No one descriptor performed the best for all kinds of retrieval tasks. Each descriptor had its own strength and weakness across retrieval tasks.

There are three broad categories of 3D object representations: feature-based methods, graph-based methods, and other representations, such as the medial axis representation. Feature-based methods, which are the most popular, can be further categorized into (1) global features, (2) global feature distributions, (3) spatial maps, and (4) local features. Local features are often the salient points of a 3D object, computed in various ways. Most methods use the curvature properties of the 3D object to find the salient points on the surface of the object [17, 7]. Castellani [5] proposed a new methodology for detection and matching of salient points based on measuring how much a vertex is displaced after filtering. The salient points are then described using a local description based on Hidden Markov Models (HMM), and HMM similarity measures are used for the 3D

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MIR'08, October 30–31, 2008, Vancouver, British Columbia, Canada.
Copyright 2008 ACM 978-1-60558-312-9/08/10 ...\$5.00.

object retrieval. Ohbuchi et al. [11] rendered multiple views of a 3D model and extracted local features from each view image using the SIFT algorithm. The local features were then integrated into a histogram using a bag-of-features approach to retrieval. Novatnack et al. [10, 9] extracted corners and edges of a 3D model by first parameterizing the surface of a 3D mesh model on a 2D map and constructing a dense surface normal map. They then constructed a discrete scale-space by convolving the normal map with Gaussian kernels of increasing standard deviation. The corners and edges detected at individual scales were then combined into a unified representation for the 3D object.

Assfalg et al. [1] captured the shape of a 3D object using the curvature map of the object’s surface. The model surface is first warped into a sphere, and curvature information is then mapped onto a 2D image using the Archimedes projection. Retrieval of 3D objects is performed by comparing the 2D map of the query object against the 2D maps of the database models using a histogram-based search and a weighted walkthrough of the map regions. Our method is quite related to that of [1], but it differs from theirs in that it does not use curvature information directly. Instead it uses a classifier to find salient points and labels them according to the classifier prediction scores. It then uses the longitude and latitude positions of salient points on the object’s surface to create a 2D signature. A classifier is trained on these 2D map signatures for classification and retrieval purposes.

3. DATA ACQUISITION

The 3D objects used in our experiments were obtained by scanning hand-made clay toys with a Roldand-LPX250 laser scanner with a maximal scanning resolution of 0.008 inches for plane scanning mode [12]. Raw data from the scanner consisted of clouds of 3D points that were further processed to obtain smooth and uniformly sampled triangular meshes of 0.9-1.0 mm resolution.

Fifteen objects were originally scanned to create a 7-class database. In this work we focus on heads of different classes and shapes to emphasize the repeatability of the method. The seven classes are: cat head, dog head, human head, rabbit head, horse head, tiger head and bear head. Each of the fifteen objects were randomly morphed using the 3D Studio Max operators for tapering, twisting, bending, stretching, and squeezing to increase the database size. A total of 250 morphed models per original object were created. For this work, we used 75 morphed models from each of the original objects for training and testing the classifier. Points on the morphed models are in full correspondence with the original models from which they were constructed.

In addition to the heads, we also performed retrieval experiments on a subset of the SHREC 2008 Classification Database.

4. LEARNING SALIENT POINTS

Our methodology starts by applying a low-level operator to every point on the surface mesh. The low-level operators extract local properties of the surface points by computing a single feature value v_i for every point p_i on the mesh surface. In this work, we will use the absolute values of the Gaussian curvature for our experiments. The low-level feature values are convolved with a Gaussian filter to reduce the noise. Figure 1a shows an example of the absolute Gaussian cur-




Figure 1: (a) Absolute Gaussian curvature low-level feature value and (b) Smooth low-level feature values after convolution with the Gaussian filter.

vature values of a 3D model and Figure 1b shows the results of applying a Gaussian filter over the low-level values.

After this first step, every point p_i on the surface mesh will have a low-level feature value v_i . The second step performs mid-level feature aggregation to compute a number of values for a given neighborhood of every point p_i on the surface mesh. In this work, we use a local histogram to aggregate the low-level feature values of each point. The histograms are computed by taking a neighborhood around each point and accumulating the low-level features in that neighborhood. The size of the neighborhood is determined by multiplying a constant c , $0 < c < 1$, with the diagonal of the object’s bounding box. This ensures that the size of the neighborhood is scaled according to the object’s size and that the results are comparable across objects. The aggregation results in a d -dimensional vector f_i for every point p_i on the surface mesh, where d is the number of histogram bins. For our experiments, we used $d = 250$ and $c = 0.05$.

Because we were not satisfied with the salient points computed by standard interest operators (eg. Kadir’s entropy operator and Lowe’s SIFT operator applied to 3D), we chose to teach a classifier the characteristics of points that we regard as salient. Histograms of low-level features are used to train a Support Vector Machine (SVM) classifier [13, 16] to learn the salient points on the 3D surface mesh. We used the SVM implemented in WEKA for our experiments [18]. The training data for supervised learning for the classifier are obtained by manually marking salient and non-salient points on the surface of each training object. An average of 12 salient points and 12 non-salient points were marked on the training objects. Since our current database contains head shapes (human heads, wildcat heads, bear heads etc), the salient points that were chosen included the tip of the nose, corners of the eyes, corners and midpoints of the lips, etc. The histogram of low-level features of each of the marked points were saved and used for the training. Figure 2 shows examples of manually marked salient and non-salient points on the training data. Since the morphed models were in total correspondence with the original model, we only had to manually mark the original models.

A small training set, consisting of 25 morphs of the cat head model, 25 morphs of the dog head model, and 50 morphs of human head models was used to train the classifier to learn the characteristics of the salient points in terms of the histograms of their low-level features. After training is complete, the classifier is able to label each of the points of any 3D object as either salient or non-salient and provides a confidence score for its decision. A threshold T is applied to the confidence scores for the salient points. In our experi-




Figure 2: Examples of manually marked salient (blue color) and non-salient (red color) points on (a) cat head model, (b) dog head model, and (c) human head model.




Figure 3: Salient point prediction for (a) cat head class, (b) dog head class, and (c) human head class. Non-salient points are colored in red, while salient points are colored in different shades ranging from green to blue, depending on the classifier confidence score assigned to the point. A threshold ($T = 0.95$) was applied to include only salient points with high confidence scores.

ments, we used $T = 0.95$ to keep only the salient points with high confidence scores from the classifier. Figure 3 shows results of the salient points predicted on instances of the cat, dog and human head class. The salient points are colored according to the classifier confidence score assigned to the point. Non-salient points are colored in red, while salient points are colored in different shades of blue with dark blue having the highest prediction score. While the classifier was only trained on cat heads, dog heads, and human heads, it does a good job of finding salient points on the other classes of heads, and the 3D patterns produced are quite repeatable across objects of the same class.

5. 2D LONGITUDE LATITUDE MAP SIGNATURE

Most 3D object classification methods require the use of a 3D descriptor or signature to describe the shape and properties of the 3D objects. Our signature is based on the salient point patterns of the 3D object mapped onto a 2D plane via a longitude-latitude transformation.

Before mapping the salient point patterns onto the 2D plane, the salient points are assigned a label according to the classifier confidence score assigned to the point. The classifier confidence score range is discretized into a number of bins. For our experiments, at confidence level 0.95 and above, we chose to discretize the confidence score range into 5 bins. Each salient point on the 3D mesh is assigned a label based on the bin into which its confidence score falls.

To obtain the 2D longitude-latitude map signature for an object, we calculate the longitude and latitude positions of




Figure 4: Salient point patterns on 3D objects of Figure 3 and their corresponding 2D longitude-latitude map signatures.

all the 3D points on the object’s surface. Given any point p_i (p_{ix}, p_{iy}, p_{iz}), the longitude position θ_i and latitude position ϕ_i of point p_i are calculated as follows:

$$\theta_i = \arctan\left(\frac{p_{iz}}{p_{ix}}\right) \quad \phi_i = \arctan\left(\frac{p_{iy}}{\sqrt{p_{ix}^2 + p_{iz}^2}}\right)$$

where $\theta_i = [-\pi, \pi]$ and $\phi_i = [-\frac{\pi}{2}, \frac{\pi}{2}]$.

A 2D map of the longitude and latitude positions of all the points on the object’s surface is created by binning the longitude and latitude values of the points into a fixed number of bins. A bin is labeled with the salient point label of the points that fall into that bin. If more than one label is mapped to a bin, the label with the highest count is used to label the bin. Figure 4 shows salient point patterns for the cat head, dog head, and human head model of Figure 3 and their corresponding 2D map signatures. Figure 5 shows how different objects that belong to the same class will have similar 2D longitude-latitude signature maps.

6. 3D OBJECT RETRIEVAL

By creating a signature for each of the 3D objects, we are able to perform similarity-based retrieval for all the objects in the database. Retrieval of 3D objects in the database is done by calculating the distance between the 2D longitude-latitude map signature of a query object and the 2D longitude-latitude map signatures of the objects in the database. We treat the 2D maps as vectors and use Euclidean distance as the distance measure.

The retrieval performance is measured using the average normalized rank of relevant images [8]. The evaluation score for a query object q is calculated as follows:

$$score(q) = \frac{1}{N \cdot N_{rel}} \left(\sum_{i=1}^{N_{rel}} R_i - \frac{N_{rel}(N_{rel} + 1)}{2} \right)$$

where N is the number of objects in the database, N_{rel} is the number of database objects that are relevant to the query object q (all objects in the database that have the same class label as the query object), and R_i is the rank assigned to the i -th relevant object. The evaluation score ranges from 0 to 1, where 0 is the best score as it indicates that all database objects that are relevant are retrieved before all other objects in the database. A score that is greater than 0 indicates that some nonrelevant objects are retrieved before all relevant




Figure 5: Objects that are similar and belong to the same class will have similar 2D longitude-latitude signature maps.

objects. Our experiments are divided into two parts: pose-normalized retrieval and rotation-invariant retrieval.

6.1 Pose-Normalized Retrieval

For the pose-normalized retrieval experiments, all objects in the database were rotated to face forward and have the same orientation. For these experiments, the database consisted of 105 objects which were categorized into 7 classes: cat head, dog head, human head, rabbit head, horse head, wildcat head, and bear head. The retrieval performance was measured over all 105 objects using each in turn as a query object. We performed two different retrieval experiments for the pose-normalized experiments on our head data.

For the first retrieval experiment, when calculating the evaluation score for a query object, the relevant objects were defined to be those objects that are morphed versions of the query object. This resulted in 15 relevant labels for the database. The algorithm was able to obtain a score of 0 for almost all the queries, except for one of the horse head queries that obtained a score of 0.00136054. For this horse query, a morphed version of a different horse head was returned at rank 6 before the last morphed version of the query horse. Since the two horses had different labels, the retrieval score was not zero.

The second retrieval experiment had a similar setup to the first experiment. The difference was in categorizing which objects were to be considered relevant to the query object. For this experiment, the relevant objects were all objects in the same general class: human, cat, dog, horse, rabbit, wildcat, and bear heads. This resulted in a total of 7 rele-

class	mean	stddev
cat head	0	0
dog head	0	0
human head	0	0
rabbit head	0	0
horse head	0.064	0.069
wildcat head	0.263	0.046
bear head	0	0

Table 1: Pose-normalized retrieval experiment 2: the mean and standard deviation of the evaluation scores for all 7 head classes. The objects in the experiments were pose normalized, and the relevant objects were all objects that belonged to the same general class.

vant labels instead of 15. All the scores are extremely low with the exception of the wildcat head. Table 1 shows the mean and standard deviation of the evaluation scores for each class, using each object in the database as a query.

We also tested the performance of our 2D longitude-latitude signature map for retrieval on a subset of the SHREC2008 Classification Database [14]. Our database consisted of 100 models pre-classified into three different levels of categorization ranging from a coarse level to a very fine level of categorization. The objects in the database were pose normalized, so that objects that belong to the same class were in the same orientation. The salient points on the SHREC 2008 database objects were identified using the same classi-

Class #	Class Name	# Objects	Mean	Stddev
1	human-diff-pose	17	0.422	0.213
2	four-legged-animal	12	0.313	0.066
3	hourglass,chess-piece	13	0.431	0.133
4	knots	8	0.021	0.008
5	plane,heli,missile	12	0.421	0.063
6	pipes,spiral	19	0.283	0.062
7	vase,bottle,teapot	19	0.252	0.029


Table 2: Pose-normalized retrieval experiment 3: the mean and standard deviation of the evaluation scores for the SHREC 2008 database with coarse categorization.

Class #	Class Name	# Objects	Mean	Stddev
1	human-dif- pose	12	0.457	0.171
2	toy-diff-pose	5	0.356	0.274
3	4-legged-long-tail	3	0.347	0.163
4	4-legged-animal	9	0.254	0.051
5	hour-glass	2	0.15	0.212
6	chess piece	2	0.05	0.071
7	statues	9	0.354	0.218
8	knots	4	0.028	0
9	torus	4	0.052	0.021
10	airplane	6	0.252	0.078
11	heli	4	0.059	0.019
12	missile	2	0.233	0.159
13	rounded-pipe-part	1	0	-
14	square-pipe-part	5	0.187	0.024
15	long-pipe-part	2	0.068	0.018
16	spiral	7	0.262	0.023
17	scissor	3	0.238	0.238
18	wheel	1	0	-
19	vase	4	0.179	0.041
20	bottle	5	0.063	0.055
21	teapot	2	0.215	0.134
22	mug	4	0.141	0.078
23	vase with handle	4	0.425	0.064

Table 3: Pose-normalized experiment 4: the mean and standard deviation of the evaluation scores for the SHREC 2008 database with mid-level categorization.

fier that was trained on the head database. No additional training on salient points was performed. The retrieval performance was measured over all 100 objects using each in turn as a query object. We calculated the retrieval score mean and standard deviation for each class for each level of categorization. The coarse categorization divided the object into 7 classes, while the mid-level categorization had 23 class labels. The finest level of categorization was too fine and resulted in many single-object classes, which were not useful when evaluating the retrieval performance. In this work, we show the retrieval performance for the coarse-level and mid-level categorization. Table 2 shows the mean and standard deviation retrieval scores for the coarse-level categorization, while Table 3 shows the mean and standard deviation retrieval scores for the mid-level categorization.

Figure 4 shows a retrieval example using a subset of the SHREC 2008 classification database. The query object is a four-legged animal. The results show that the top four retrieved objects also belong to the four-legged-animal class, and that the signatures obtained for these animals are quite similar. The next retrieved object at rank 5 is a pipe part whose orientation is similar to these four-legged animals and

Query name	Query object	Query signature
24		

Retrieval results:









Rank	Name	3D Object	Signature	Dist
1	23			54.09
2	25			55.34
3	22			55.8
4	28			57.45
5	70			60.83
...
100	13			178.54

Table 4: Retrieval example using a subset of the SHREC 2008 Classification Database. The query example is one of the four-legged animal. The top four retrieved object are also four-legged animals.

had salient points detected in similar locations (front and back of the pipe) as the four-legged animals.

6.2 Rotation-Invariant Retrieval

Most 3D objects created for databases and applications are not pose normalized; hence it is essential for our approach to be rotation invariant. The 2D signatures we generate are dependent on the location and orientation of a 3D coordinate system associated with the model. To achieve rotation invariance, we created multiple rotations of each 3D

object from which we generated new 2D longitude-latitude signatures. We showed in prior work [2] on object classification that classification accuracy decreased as a function of number of rotations per objects in the training set. For the head database, we found that accuracy was above 90% for rotations generated at 100 degree increments. Therefore, in this paper, the rotated versions were generated at 100 degree increments for all three axes resulting in $4 \times 4 \times 4$ rotated signatures for seven morphs of each of the fifteen original object in the database. A total of 6720 rotated signatures were stored in the database. Retrieval was performed by calculating the distance between the query object signature and all the rotated signatures in the database. The retrieval algorithm then selected the best rotated signature for each original object and returned the object at that rotation. Hence, each query retrieval had $15 \times 7 = 105$ objects returned in order of similarity.

We ran two experiments to test the rotation invariance properties of the retrieval using randomly-rotated objects as the queries. These rotations could be at any angles ($\theta_x, \theta_y, \theta_z$) about the 3 axes, not just at multiples of 100 degrees, so the maps could be quite different from those in the database. We generated 10 rotations per query object resulting in 1050 queries in total. The experiments differed in how the relevant objects were categorized. The first experiment evaluated relevant objects that were a morphing version of the query object, while the second experiment used the general class label for the relevant object’s label. Table 5 and Table 6 shows the mean and standard deviation of the evaluation scores for these two experiments. For the first experiment, most of the non-relevant retrievals were actually morphed versions of objects in the same general class. For example, when querying on human head 1, some morphed versions of human head 2 and human head 3 were retrieved before the rest of human head 1. Similar scenarios happened for the leopard, tiger and lion head queries. When querying with a leopard head, some morphed versions of the tiger and lion heads were returned before all relevant leopard head models were returned. In the second experiment, when querying with a leopard head, some bear heads were retrieved before all leopard heads had been retrieved.

Table 7 shows an example of a retrieval. The query for the retrieval is a human head that is randomly rotated at rotation angle (188,164,139). The evaluation score for this query was 0.276 in the first experiment and 0.019 in the second experiment. The results show that all the human heads in the database were retrieved before the rest of the objects. Table 8 shows another good retrieval example. The query for the retrieval is a rabbit head that is randomly rotated at angle (345,304,278). The evaluation score for this query was 0 in both experiments. Note that the retrieval results may be at a different rotation than the query, since we only stored rotations in 100 degree increments about each axis.

Table 9 shows an example of incorrect retrieval results. The query for this retrieval was a leopard head that was randomly rotated with angles (64,4,146). The evaluation score for this retrieval was 0.571429 in the first experiment and 0.64723 in the second experiment. The results show that the first object retrieved was actually a bear head, and that the first leopard head object in the database was retrieved at rank 29.

class label	mean	stdev
cat	0.274	0.151
dog	0.345	0.307
human-1	0.192	0.114
human-2	0.1677	0.109
human-3	0.0619	0.063
human-4	0.058	0.045
human-5	0.072	0.049
rabbit-1	0.249	0.253
rabbit-2	0.165	0.165
horse-1	0.242	0.168
horse-3	0.087	0.137
leopard	0.268	0.280
tiger	0.165	0.216
lion	0.093	0.125
bear	0.204	0.188

Table 5: Rotation-Invariant Retrieval Experiment 1: The mean and standard deviation of the evaluation scores for all 15 classes. The query objects for the retrieval were randomly rotated and the relevant objects were all morphed versions of the query object.

class label	mean	stdev
cat	0.274	0.151
dog	0.345	0.307
human	0.009	0.013
rabbit	0.197	0.156
horse	0.326	0.187
wildcats	0.258	0.208
bear	0.201	0.188

Table 6: Rotation-Invariant Retrieval Experiment 2: The mean and standard deviation of the evaluation scores for all 7 head classes. The query objects for the retrieval were randomly rotated and the relevant objects were based on general classes.

7. CONCLUSIONS AND FUTURE WORK

We have described a learning approach to 3D object retrieval. Supervised learning is performed on selected salient points on the training data. The classifier learns the characteristics of the salient points and is able to predict salient points on new objects. The patterns of the salient points are used to train a second classifier by transforming the patterns to a 2D map using the longitude-latitude transformation to produce a 3D object signature. Using a simple Euclidean distance function between two 2D map signatures, we are able to perform similarity-based retrieval for 3D objects. Experimental results show that the signature can be used to retrieve similar objects to a query both for objects that are pose normalized and those that are rotated randomly.

We are investigating clustering the rotated signatures in the database to further reduce space usage and computation time so that we can add more objects and classes. We are also working on the use of a multi-class classifier as a first step before performing retrieval. Given a query object, the trained classifier is used to predict the best class labels for the query. Distance computation is then only performed between the query and database objects that are





Query name	Query object	Query signature	Retrieval name	Retrieval Object	Retrieval signature	Dist
leopard_head_30 rotated(64,4,146)			bear_head_2_28 rotated(45,0,45)			100.4

Table 9: Example of incorrect retrieval results. The query for the retrieval is a leopard head rotated at angle (64,4,146). The top retrieved object is a bear head, which does show similarity to the leopard head.



in the predicted classes. In addition, we are working on finding a mapping function between two longitude-latitude signatures. This will eliminate the need to generate rotated versions of the database object signatures and will allow us to handle much larger databases.

8. ACKNOWLEDGMENTS

This research is supported by the National Science Foundation under grant number DBI-0543631.

9. REFERENCES

- [1] J. Assfalg, A. D. Bimbo, and P. Pala. Content-based retrieval of 3D models through curvature maps: a cbr approach exploiting media conversion. *Multimedia Tools Appl.*, 31:29–50, 2006.
- [2] I. Atmosukarto and L. G. Shapiro. A learning approach to 3D object representation for classification. *Submitted to. S+SSPR*, 2008.
- [3] A. D. Bimbo and P. Pala. Content-based retrieval of 3D models. *ACM Trans. on Multimedia Computing, Communications, and Applications*, 2(1):20–43, 2006.
- [4] B. Buston, D. Keim, D. Saupe, and T. Schreck. Content-based 3D object retrieval. *Computer Graphics and Applications*, 27(4):22–27, 2007.
- [5] U. Castellani, M. Cristani, S. Fantoni, and V. Murino. Sparse points matching by combining 3D mesh saliency with statistical descriptors. *Computer Graphics Forum*, 27(2):643–652(10), 2008.
- [6] N. Iyer, S. Jayanti, K. Lou, Y. Kalyanaraman, and K. Ramani. Three-dimensional shape searching: state-of-the-art review and future trends. *Computer Aided Design*, 37:509–530, 2005.
- [7] C. H. Lee, A. Varshney, and D. W. Jacobs. Mesh saliency. *ACM Trans. Graph*, 24(3):659–666, 2005.
- [8] H. Müller, S. Marchand-Maillet, and T. Pun. The truth about corel - evaluation in image retrieval. In *CIVR*, 2002.
- [9] J. Novatnack and K. Nishino. Scale-dependent 3D geometric features. In *Proc. International Conference on Computer Vision (ICCV)*, 2007.
- [10] J. Novatnack, K. Nishino, and A. Shokoufandeh. Extracting 3D shape features in discrete scale-space. In *Proc. 3DPVT*, 2006.
- [11] R. Ohbuchi, K. Osada, T. Furuya, and T. Banno. Salient local visual features for shape-based 3D model retrieval. In *Shape Modeling International*, 2008.
- [12] S. Ruiz-Correa, L. G. Shapiro, M. Meila, G. Berson, M. L. Cunningham, and R. W. Sze. Symbolic signatures for deformable shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(1):75–90, 2006.
- [13] B. Scholkopf and A. J. Smola. *Learning with kernels*. Cambridge University Press, 2002.
- [14] SHREC2008. Classification of watertight models track http://shrec.ge.imati.cnr.it/shrec08_classification.html.
- [15] J. Tangelder and R. Veltkamp. A survey of content-based 3D shape retrieval methods. In *Shape Modeling International*, 2004.
- [16] V. V. Vapnik. *Statistical Learning Theory*. John Wiley and Sons, 1998.
- [17] K. Watanabe and A. G. Belyaev. Detection of salient curvature features on polygonal surfaces. *Comput. Graph. Forum*, 20(3), 2001.
- [18] I. H. Witten and E. Frank. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann San Fransisco, 2nd edition, 2005.

Query name	Query object	Query signature
human_1_29 rotated (188,164,139)		

Retrieval results:





Info	Name	3D Object	Signature
Rank=1 Dist=77.91	human_3_26 rotated $\theta_x=0$ $\theta_y=0$ $\theta_z=0$		
Rank=2 Dist=78.15	human_3_31 rotated $\theta_x=0$ $\theta_y=0$ $\theta_z=0$		
Rank=3 Dist=78.26	human_3_27 rotated $\theta_x=0$ $\theta_y=0$ $\theta_z=0$		
...
Rank=29 Dist=89.03	horse_1_28 rotated $\theta_x=0$ $\theta_y=0$ $\theta_z=100$		
...
Rank=105 Dist=115.97	wildcat_27 rotated $\theta_x=100$ $\theta_y=300$ $\theta_z=300$		

Table 7: Retrieval example using human head rotated at angle (188,164,139) as the query. The human heads in the database were all retrieved before other objects in the database.

Query name	Query object	Query signature
rabbit_1_27 rotated (345,304,278)		

Retrieval results:








Info	Name	3D Object	Signature
Rank=1 Dist=84.33	rabbit_1_30 rotated $\theta_x=0$ $\theta_y=300$ $\theta_z=300$		
Rank=2 Dist=84.34	rabbit_1_27 rotated $\theta_x=0$ $\theta_y=300$ $\theta_z=300$		
Rank=3 Dist=84.59	rabbit_1_26 rotated $\theta_x=0$ $\theta_y=300$ $\theta_z=300$		
...
Rank=15 Dist=94.91	horse_1_28 rotated $\theta_x=0$ $\theta_y=300$ $\theta_z=300$		
...
Rank=105 Dist=114.9	lion_31 rotated $\theta_x=0$ $\theta_y=300$ $\theta_z=300$		

Table 8: Retrieval example using rabbit head rotated at angle (345,304,278) as the query. The rabbit heads in the database were all retrieved before other objects in the database.