Invariant Local Shape Descriptors: Classification of Large-Scale Shapes with Local Dissimilarities

Xizhi Li
University of Bremen
lixizhi@uni-bremen.de

Patrick Lange
University of Bremen
lange@cs.uni-bremen.de

René Weller
University of Bremen
weller@cs.uni-bremen.de

Gabriel Zachmann
University of Bremen
zach@cs.uni-bremen.de

ABSTRACT

We present a novel statistical shape descriptor for arbitrary three-dimensional shapes as a six-dimensional feature for generic classification purposes. Our feature parameterizes the complete geometrical relation of the global shape and additionally considers local dissimilarities while being invariant to the shape appearance. Our approach allows the classification of large-scale shapes with only small local dissimilarities. Our feature can be easily quantized and mapped into a histogram, which can be used for efficient and effective classification. We take advantage of GPU processing in order to efficiently compute our invariant local shape descriptor feature even for large-scale shapes. Our synthetic benchmarks show that our approach outperforms state-of-the-art methods for local shape dissimilarity classification. In general, it yields robust and promising recognition rates even for noisy data.

CCS CONCEPTS

• **Information systems → Similarity measures**; • **Computing methodologies → Concurrent programming languages**;

KEYWORDS

Shape descriptor, Surface representation, Large-scale models

1 INTRODUCTION

Robust scene interpretation in the form of detection and classification of previously known 3D objects in arbitrary scenes is a key factor in various computer vision approaches. Efficient and smart shape descriptors are fundamental to object detection and classification. According to Wahl et al. [25] and Rusu et al. [21] such shape representation have to be

(1) compact,
(2) robust,
(3) they should be invariant, e.g. not depending on a global coordinate frame, and they
(4) should have the descriptive capacity to distinguish arbitrary shapes.

Often, e.g. in robotics, the object detection and classification is done on point cloud streams. Due to a consecutively improving 3D sensing technology (stereo systems, laser scanner, or consumer electronics such as the Kinect) these point clouds do not only become larger but additionally contain more details. Consequently, it is necessary that object detection and classification approaches adapt to the increased (local) 3D object detail. Hence, we identified the following additional challenges for such kinds of shape descriptors:

(5) they should consider small local dissimilarities,
(6) their computation should be manageable to handle also large-scale shapes.

Wahl et al.'s work effectively solve the first four items by creating a surflet-pair histogram to represent the shape of 3D objects and matching histograms with KL divergence [10]. However, their approach fails to deal with the small locally dissimilar and large-scale computing problems. The main reason for the performance issues are the sequential computation of the histograms and the usage of the statistical KL method for the classification. On the other hand, Zhang et al. [26] already showed that the GPU can be applied to accelerate the 3D object retrieval process. Moreover, machine learning algorithms have become a very popular and powerful tool for 2D or 3D object classification and regression problems, especially for large, high-dimension histograms. In this paper we present a substantial extension to the approach by Wahl. In detail, our contributions are:

- a novel *local* feature that considers small local characteristics of the object,
- a parallelization of the histogram computation,
- a machine-learning-based classification algorithm that can handle large-scale shapes.

Basically, we augment the the four-dimensional feature of Wahl et al. into a six-dimensional feature that considers additionally the local Gaussian curvature and the local angle of the object. Our approach is invariant to translation, rotation and scale of the shapes and moreover, it is robust to noise. We have implemented our algorithm using CUDA that allows it to completely run on the GPU. As a use case scenario we chose the classification of 3D asteroids from point cloud data. This scenario is typical for large-scale objects with local dissimilarities and is currently discovered in spacecraft operation studies for autonomous landing [19]. Additionally, we evaluated regular objects from the NTU database. We compared our

CGI '17, June 27-30, 2017, Yokohama, Japan

ACM ISBN 978-1-4503-5228-4/17/06…$15.00
https://doi.org/10.1145/3095140.3095149

approach to several state-of-the-art shape descriptors. Our results show that our approach is capable of classifying both large-scale shapes with local dissimilarities based on their local appearance and standard shapes based on their global and local dissimilarities efficiently.

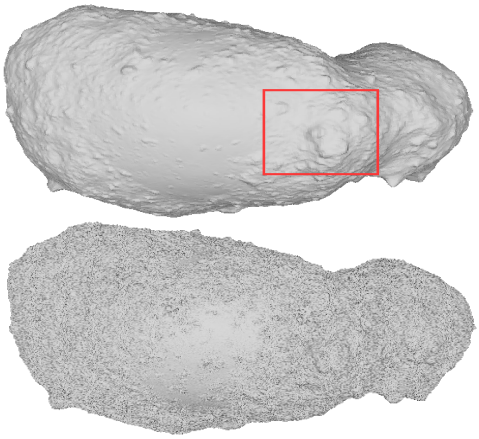

Figure 1: The dissimilarity of different resolution 3D models of Itokawa.

2 RELATED WORK

The goal of shape-based object classification is to formulate shape properties which accurately represent the object. This shape descriptors are used to efficiently classify them while focusing on (1) the compactness of shape descriptors and (2) the robustness of shape descriptors. A detailed overview of shape descriptors for shape-based object classification can be found in [15, 24]. Here, we give an overview of relevant approaches which are directly related to our work.

In early research, Bay [2] established the SURF detector, and [17] proposed SIFT using local invariant descriptors. More recent approaches, e.g. from [3, 5, 8, 16, 25] focused on new shape descriptors. Namely, [25] introduced a four- dimensional global shape descriptor. They defined a 3D coordinate frame for each pair of oriented points (so-called surflet-pairs), and defined a four-dimensional, pose invariant shape descriptor, which describes these surflets. [3] proposed the view-based global shape descriptors Light Field descriptor(LFD) which aims at describing 3D models by a set of two-dimensional representations. In contrast, [8] introduced two-dimensional spherical harmonics based shape descriptors. This approach does not contain a sophisticated classification scheme because the similarity between two shapes is calculated by the Euclidean between two spherical harmonic descriptors. Lo and Siebert [16] proposed Trift which extended the idea of SIFT from 2D image to 2.5D domain. Their idea is to fusion the histogram of range surface topology types with the histogram of the range gradient orientations to form a new feature descriptor.

Another spectral-based shape analysis method called shapeDNA [20] was proposed by Reuter. Their method extracts fingerprints of an arbitrary surface by taking the eigenvalues of its respective Laplace-Beltrami operator. This is the basis for a series of shape descriptors based on such Laplace-Beltrami opearators, such as the wave kernel signature (WKS)[1] or the scale-invariant heat kernel signature (SIHKS)[12].

The approaches mentioned above trace the development on 3D shape analysis from early general shape description to recent spectral shape analysis. However, none of them considers local shape dissimilarities of large-scale objects and and moreover, they are susceptible to noise that often appears in point clouds.

3 OUR DESCRIPTORS

Wahl et al.'s four-dimensional geometric feature is the basis of our novel six-dimensional geometric feature. Basically, they proposed surflet-pair histograms to describe the shape of 3D objects. We start with a short recap of this approach.

3.1 Recap: Surflet-Pair-Relation Histograms

An important advantage of Wahl et al.'s approach is its transformational invariance. In order to realize this, he introduced a canonical coordinate system by extracting features U, V, W (see Figure 2(a)) as a transformation independent reference. They defined a canonical coordinate system as follows:

1 Extract the whole pairwise points and its normals from the surface mesh of object (P_i, n_i). Randomly select surflet-pairs (P_i, n_i) and (P_j, n_j). If satisfy P_i satisfies Equation 1 we simply set P_i as the origin of the canonical coordinate system otherwise P_j.

$$\left|n_j \times (p_j - p_i)\right| \leqslant \left|n_j \times (p_j - p_i)\right| \tag{1}$$

2 Then he constructs the canonical coordinate system by computing U, V, W as the base vectors: Assuming P_i as the origin so n_i is U, we normalize the vector $P_i - P_j$ by $\rho = \frac{P_i - P_j}{|P_i - P_j|}$ in order ensure that the feature is invariant to scaling. The canonical coordinate system is then given by $V = U \times \rho, W = U \times V$.

From this canonical coordinate system, we derive Wahl et al.'s global features as follows: Given an object represented by a point set P, for each pair of points we define four features for the complete four-dimensional vector G for each point P_i:

$$G\,(\vec{P_i}, \vec{P_j}) = (\alpha, \beta, \gamma, \delta) \tag{2}$$

by

- $\alpha = \arctan(\vec{W} \times \vec{n_2}, \vec{U} \times \vec{n_2}), \alpha \in -\frac{\pi}{2}, \frac{\pi}{2},$
- $\beta = \vec{V} \times \vec{n_2}, \beta \in (-\pi, \pi)$
- $\gamma = \vec{U} \times \vec{\rho}, \gamma \in (-\pi, \pi)$
- $\delta = \frac{|P_1 - P_2|}{Max(|\vec{P_i} - \vec{P_j}|)}, \delta \in (0, 1).$

3.2 Our Adaptive Hybrid Shape Descriptor

Our novel invariant local geometric features extend the basic four-dimensional $G(\vec{P_i})$ by two additional dimensions. They are inspired by the human cognition. Early psychophysical experiments showed that human visual system decomposes complex shapes into parts based on curvature and processes salient features before higher level recognition [23]. This research motivates us to focus on curvature to represent the local shape of the 3D model. This local geometric feature $L(P_i)$ for each point can be represented by the following

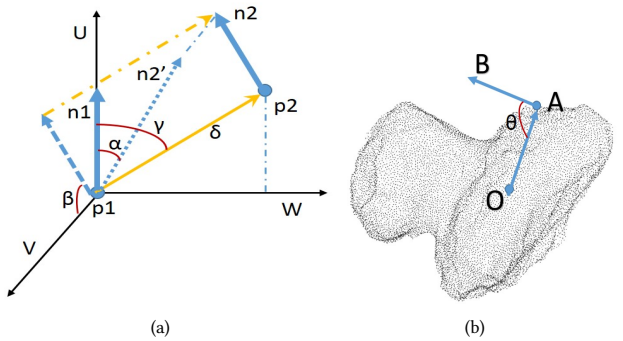

Figure 2: (a) Pairwise points P_1,P_2 and their normal vectors n_1, n_2. U, V, W are the intrinsic reference frame built by the pairwise points. And n_2'fi is the projection of n_2 to UW plane, α , β , γ represents the angles between (n_1, $n_2$fi) , (V, n_2), and (n_1, $P_2 - P_1$). δ is the distance between the pairwise points.(b) O denotes the geometric center of the object. Vector OA and AB represents the position vector and normal vector of vertex A, θ is the inclined angle.

two parameters η and κ:

$$L(\vec{P}_i) = (\eta, \kappa) \tag{3}$$

with

- η is the Gaussian curvature
- κ represents the local normal perturbation of the object's normal.

In order to obtain a smooth curvature at each point P_i, we apply the discrete curvature analysis according to [6]. Additionally, we consider the normal perturbation with the parameter κ. The main challenge is to represent the point's normal in a transformational invariant way. Let c be the geometric center of the model and n_i the normal of point P_i. Then we define κ= (P_i-c)n_i, i.e. κ represents the angle between the vector that is spanned by the object's center and the point's position and the normal of the point (see Figure 2(b)).

The global shape descriptor according to Wahl et al. and our local shape descriptor can be easily combined to our new adaptive hybrid shape descriptor (AHD):

$$AHD(\vec{P}_i, \vec{P}_j) = (G(\vec{P}_i, \vec{P}_j), L(\vec{P}_i)) = (\alpha, \beta, \gamma, \delta, \eta, \kappa) \tag{4}$$

Since G($\vec{P}_i, \vec{P}_j$) as well as L($\vec{P}_i$) are transformation invariant, also G($\vec{P}_i, \vec{P}_j$) is transformation invariant. For a point cloud consisting of n points, G($\vec{P}_i, \vec{P}_j$) can be computed in O(n) and L($\vec{P}_i$) in constant time for each individual point P_i.

4 TRAINING AND CLASSIFICATION

Our geometric feature described above is the basis for object recognition tasks. To do that, we create a database of histograms for a set of point clouds. The histograms are generated for each point cloud individually by computing our AHD-feature for all pairs of points and then discretizing them into bins. Similar to Wahl et al., we chose five bins per dimension. This results in a total number of $5^6 = 15625$ bins for each object. In the end, we get a 15625-dimensional vector that represents the object.

4.1 Parallelization

A nice property of our geometric feature is that the histogram generation can be easily parallelized. Obviously, the parameters for each pair of points can be computed independently (see Algorithm 2). In order to bin the resulting six-dimensional vectors, we additionally have to sort these vectors. In detail, we use a parallel bitonic sort and a parallel reduction algorithm to count the number of entries per bin (see Algorithm 1). Please note, that our *local* features η and κ have to be computed only once per point, whereas the global features $\alpha, \beta, \gamma, \delta$ are computed per pair of points. Consequently, the total parallel running time of our algorithm is in O(n) assuming a perfect PRAM.

Algorithm 1: computeHistogram(Pointcloud A)

In Parallel forall *points* $p_i \in A$ **do**
    featureSet[i]=computeFeature(p_i, A)
In Parallel sort(featureSet)
histogram = **In Parallel** reduction(featureSet)

Algorithm 2: computeFeature(Point p, Poincloud A)

compute curvature η(p)
compute angle κ(p)
In Parallel forall *points* $p_i \in A$ **do**
    compute α, β, γ and δ(p, p_i)

4.2 Histogram cluster analysis

Choosing the best classification algorithm is a non-trivial task. For instance, it highly depends on the dataset and the number but also the identifiability of the clusters. For our use case of asteroid classification (see Section 5), we first tried to use the linear discriminant analysis in combination with a PCA. The results show that that the distribution can be hardly linearly divided into meaningful clusters (see Figure 3(a)). On the other hand, random forest have shown to achieve high accuracy for the classification of non-linear datasets and they can easily handle multi-class classification challenges [14]. Consequently, we decided to use random forests.

5 USE CASE: ASTEROID CLASSIFICATION

As one challenging example for the application of our algorithm we outline celestial bodies, especially asteroids. Asteroids differ in many ways from other (human created) objects because of their complex shapes, internal structures and material properties. For instance, Itokawa has significant porosities which are a key evidence for its belongingness of its corresponding taxonomic class [11]. Therefore, this kind of local dissimilarity pose a competitive challenge to our shape descriptors.

There is an increasing interest in the field of spacecraft flight to perform autonomous surface analysis and safe landing operations [19]. For these autonomous systems it is crucial to efficiently

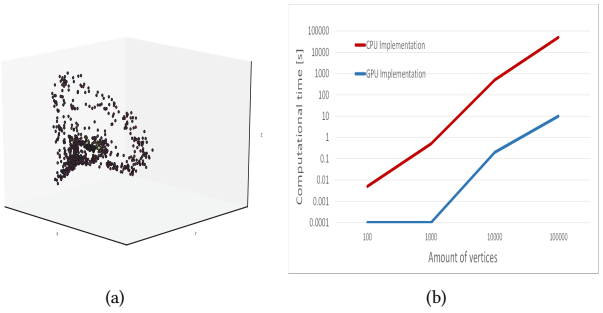

(a) (b)

Figure 3: (a) Using PCA transform to efficiently reduce raw feature histogram from 15625 dimensions to 3 dimensions. The asteroid classes are color-coded. (b) Computation times of CPU VS GPU

and accurately classify and recognize the shape and local surface details for the landing operations. Thus, the need for the ability of recognizing asteroids in arbitrary scenes based on 3D point clouds without large databases has occupied an important position.

A major drawback in the classification of asteroids is the lack of data. High-quality models of asteroids are usually recorded during rare and costly spacecraft fly-by or rendezvous missions [22]. Consequently, it is hard to obtain a large database for training and classification purposes. In order to overcome this limitation, we decided to use Poisson disk sampling on an asteroid database to increase the number of available data. Poisson disk sampling is one of the most classical methods for the fast resampling of surface points [7] and it is proven to be very robust Corsini [4].

Actually, the Poisson-sampled data draws a special challenge to our algorithm because the generated data is usually a lower-resolution model of the original high-resolution asteroid. Hence, it may lack some details on the surface. Figure 4 illustrates the Poisson sampled asteroid models.

6 EVALUATION

We have implemented our adaptive hybrid shape descriptor (AHD) in Python 3.5. We performed our experiments on a machine with Intel Core i7 quad core processor with Hyperthreading enabled, 8 GB of memory, and a Nvidia Geforce GT 640M, operated by Windows 10. We applied three experiments to measure the performance as well as the quality of our shape descriptor approach.

First, we performed a comparison of the sequential CPU algorithm and our massively parallel GPU implementation for the histogram generation and the hybrid-feature computation. Second, we evaluated the quality of our approach and its competitors for our previously outlined use case study of asteroid classification for autonomous spaceflight operations. Third, we evaluated the quality of our approach and its competitors based on the standard shape NTU database set.

We compared the quality to three state-of-the-art methods, namely the 3D Harmonics [9], LightField descriptor [3] and shapeDNA [20]. Here, we used freely available open source implementations. Additionally, we compared our approach to Wahl et al.'s original implementation (Global shape descriptor). In order to find the best parameters for our random forest for this competitive evaluation, we used grid search and selected appropriate parameters for estimators, depth, lea size and split criterion.

For our quality evaluation we use the well-known precision and recall diagram. Each of our evaluation plots precision versus recall averaged over all classified models in the database. The plot axes can be interpreted as follows [8]: For each target model in class C and any number K of top matches, "recall" represents the ratio of models in class C returned within the top K matches, while "precision" indicates the ratio of the top K matches that are members of class C. A perfect retrieval result would produce a horizontal line along the top of the plot, indicating that all the models within the target object's class are returned as the top hits. Otherwise, plots that appear shifted up and to the right generally indicate superior retrieval results.

6.1 GPU-based Histogram Generation

We compared the performance of a traditional sequential CPU implementation and massively parallel GPU implementation for our histogram generation (see Figure 3(b)). Here, we used Python 3.5. and pycuda for the implementation respectively. Our first evaluation shows that the massively parallel GPU implementation easily outperforms the traditional sequential one with an increasing number of vertices. Our GPU-based implementation gradually outperforms the sequential CPU implementation by up to a factor of 1000. The GPU timings do not include transferring data between the host CPU memory to the GPU's global memory.

6.2 Asteroid Classification Study

In the second evaluation study, we evaluated our approach and its competitors for our previously outlined use case study of asteroid classification for autonomous spaceflight operations.

We randomly selected 20 asteroids from the Planetary Data System [18] and utilized our Poisson sampling approach in order to obtain a large set of asteroids 1000 for training, testing, and evaluation purposes (see Figure 4). We add some random noise to all asteroid meshes during evaluation process to simulate realistic situation in space exploration.

Our evaluation shows that our shape descriptor approach with random forest based classification outperforms the competing methods (see Figure 5). This means that our approach is the most discriminative and effective method among all evaluated approaches. Compared with the 3D harmonic descriptor, lightfield, and shapeDNA, our method owns an average of more than 70% precision rate when average the recall axis.

Even more, our methods works well for almost all classes of asteroids, except for some difficult shapes as illustrated in the confusion matrix (see Figure 6).

In summary, these good results demonstrate that, although we merely sampled 10~20% vertices from the raw meshes, our shape descriptor is able to robustly represent the noise asteroid shape while achieving high classification rate.

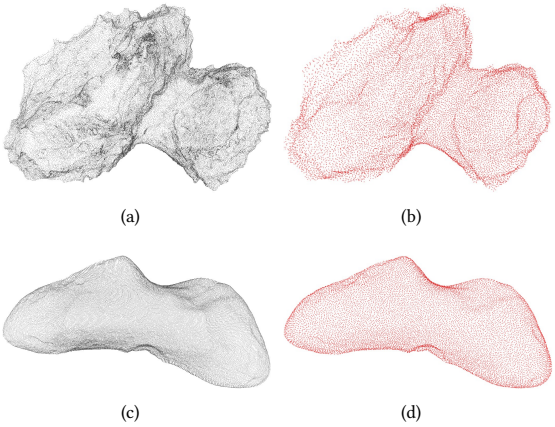

Figure 4: Example 3D asteroids and sampled asteroids. a, c are the raw asteroid named Churyumov (128,002 points) and Eros (99,846 points). b, d represent the poisson-disk sample asteroids each with 25994 and 18172 points.

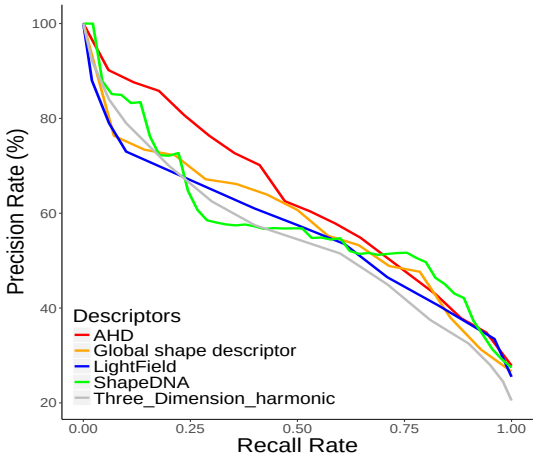

Figure 5: Plots the precision-curve for the experiment of asteroid classification.

6.3 Standard Dataset Testing

In the last evaluation study, we evaluated our approach and its competitors based on the standard shape NTU database set[1] [3]. The NTU database currently contains 10,911 3D models from 352 categories, from which we selected 1,218 representative models from the database as our testing database. These 1,218 3D models are composed into 10 classes which have the most models in each class in the database (see Table 1). Several examples of 3D models contained in these 10 most well-annotated classes are shown in Figure 7.

We split the above determined 10-class dataset randomly into training, validation and test set and used this data as the evaluation baseline. Our approach outperforms its competitors also in this evaluation study (see Figure 8). Our shape descriptor achieves the best performance with classification accuracy of 62.5% and 57.1% under invariant descriptor and global descriptor respectively, after averaging over all the recall axis. While 3D harmonic, lightfield and shapeDNA descriptors achieved 52.8%, 53.5% and 53.1% accuracy under the same conditions. Surprisingly, shapeDNA performed worst. The reason for this could be low quality and incompleteness of some meshes in the NTU database. As a result, shapeDNA is not robust enough to distinguish them. In this evaluation study, our approach does not outperform its competitors to the same extent as in the previous evaluation. We believe that the shapes of the NTU database have less local shape information than the asteroid shapes of our use-case study.

[1] http://3d.csie.ntu.edu.tw/ dynamic/database/

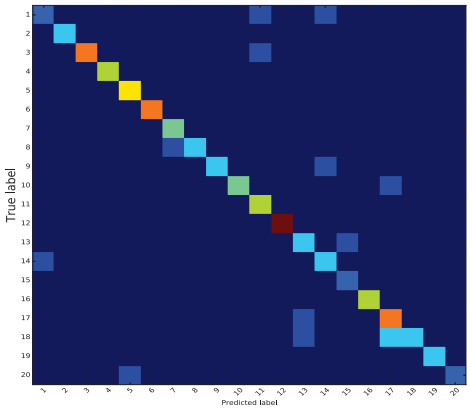

Figure 6: Plots the confusion matrix for the experiment of asteroid classification.

Table 1: Subset of NTU database

Category	Number of models	Training set
Tree	120	84
Gun	120	84
Enterprise	80	56
Wheel	78	55
Table	115	81
Potted-plant	84	60
Human	192	133
Helicopter	98	68
Fighter-plane	234	164
Four-legged-chair	97	68

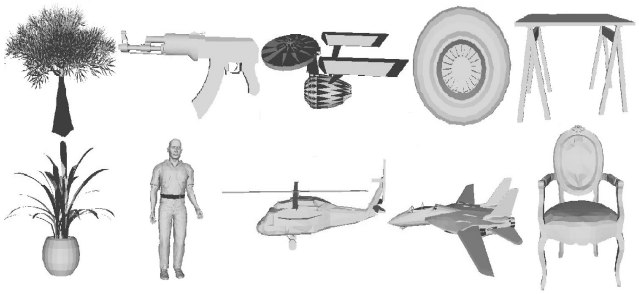

Figure 7: The example of all selected class of 3D models in NTU database.

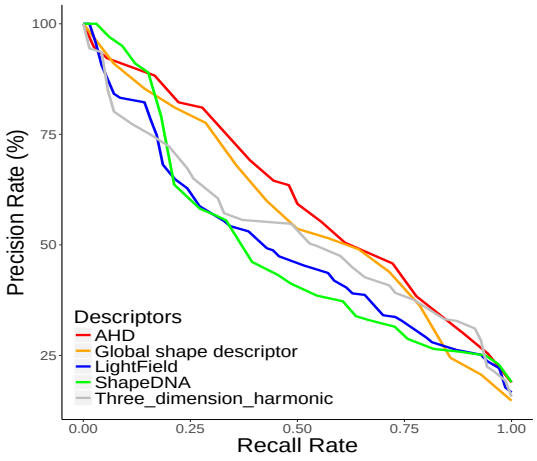

Figure 8: Precision-recall curve performance evaluation of our approach (under three schemes), compared with 3D harmonic and light field descriptors.

7 CONCLUSIONS AND FUTURE WORK

We have presented a novel statistically invariant shape descriptor for large-scale shapes with local dissimilarities. The main idea is to combine features that describe the global shape with two novel features that represent the local curvature and the normal perturbation, respectively. This enables our hybrid-feature to classify both large-scale shapes with local dissimilarities based on their local appearance and standard shapes based on their global and local dissimilarities. Our novel features are robust to noise and invariant to translation, rotation and scale of the shapes. Furthermore, we presented a parallelization of the histogram computation using GPU processing in order to deal with massive data from high-resolution 3D shapes. The results show that our GPU implementation is more than three orders of magnitude faster than the equivalent CPU implementation.

Due to its generality, our approach is applicable to a wide variety of classification domains for three-dimensional shapes. The results from our benchmarks show that our approach is able to efficiently classify large-scale shapes with local dissimilarities in a special asteroid use case but also for common objects.

In the future, we would like to further evaluate our approach with more shape databases, especially in a terrestrial context. However, we are mainly interested in improving our current approach for the outlined asteroid classification use case study. Here, we would like to incorporate reinforcement learning with our hybrid shape descriptor.

Hinton and Krizhevsky [13] proposed unsupervised deep learning method for image retrieval, this method can be a good example for our classification algorithms. Another interesting idea would be to extend our approach with additional shape descriptors to further improve it's accuracy.

REFERENCES

[1] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. 2011. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*. IEEE, 1626–1633.

[2] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. 2008. Speeded-up robust features (SURF). *Computer vision and image understanding* 110, 3 (2008), 346–359.

[3] Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. 2003. On visual similarity based 3D model retrieval. In *Computer graphics forum*, Vol. 22. Wiley Online Library, 223–232.

[4] Massimiliano Corsini, Paolo Cignoni, and Roberto Scopigno. 2012. Efficient and flexible sampling with blue noise properties of triangular meshes. *IEEE Transactions on Visualization and Computer Graphics* 18, 6 (2012), 914–924.

[5] Tal Darom and Yosi Keller. 2012. Scale-invariant features for 3-D mesh models. *IEEE Transactions on Image Processing* 21, 5 (2012), 2758–2769.

[6] Nira Dyn, Kai Hormann, Sun-Jeong Kim, and David Levin. 2001. Optimizing 3D triangulations using discrete curvature analysis. *Mathematical methods for curves and surfaces* 38, 8 (2001), 135–146.

[7] Mohamed S Ebeida, Scott A Mitchell, Andrew A Davidson, Anjul Patney, Patrick M Knupp, and John D Owens. 2011. Efficient and good Delaunay meshes from random points. *Computer-Aided Design* 43, 11 (2011), 1506–1515.

[8] Thomas Funkhouser, Michael Kazhdan, Philip Shilane, Patrick Min, William Kiefer, Ayellet Tal, Szymon Rusinkiewicz, and David Dobkin. 2004. Modeling by example. In *ACM Transactions on Graphics (TOG)*, Vol. 23. ACM, 652–663.

[9] Thomas Funkhouser, Patrick Min, Michael Kazhdan, Joyce Chen, Alex Halderman, David Dobkin, and David Jacobs. 2003. A search engine for 3D models. *ACM Transactions on Graphics (TOG)* 22, 1 (2003), 83–105.

[10] John R Hershey and Peder A Olsen. 2007. Approximating the Kullback Leibler divergence between Gaussian mixture models. In *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, Vol. 4. IEEE, IV–317.

[11] Martin Jutzi, Keith Holsapple, Kai Wünneman, and Patrick Michel. 2015. Modeling asteroid collisions and impact processes. *arXiv preprint arXiv:1502.01844* (2015).

[12] Iasonas Kokkinos, Michael Bronstein, and Alan Yuille. 2012. *Dense scale invariant descriptors for images and surfaces.* Ph.D. Dissertation. INRIA.

[13] Alex Krizhevsky and Geoffrey E Hinton. 2011. Using very deep autoencoders for content-based image retrieval.. In *ESANN*.

[14] Victor Lempitsky, Michael Verhoek, J Alison Noble, and Andrew Blake. 2009. Random forest classification for automatic delineation of myocardium in real-time 3D echocardiography. In *International Conference on Functional Imaging and Modeling of the Heart*. Springer, 447–456.

[15] Chunyuan Li and A Ben Hamza. 2014. Spatially aggregating spectral descriptors for nonrigid 3D shape retrieval: a comparative survey. *Multimedia Systems* 20, 3 (2014), 253–281.

[16] Tsz-Wai Rachel Lo and J Paul Siebert. 2009. Local feature extraction and matching on range images: 2.5 D SIFT. *Computer Vision and Image Understanding* 113, 12 (2009), 1235–1250.

[17] David G Lowe. 1999. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, Vol. 2. Ieee, 1150–1157.

[18] Susan K McMahon. 1996. Overview of the planetary data system. *Planetary and Space Science* 44, 1 (1996), 3–12.

[19] Alena Probst, Graciela Peytavi, David Nakath, Anne Schattel, Carsten Rachuy, Patrick Lange, and others. 2015. Kanaria: Identifying the Challenges for Cognitive Autonomous Navigation and Guidance for Missions to Small Planetary Bodies. In *International Astronautical Congress (IAC)*.

[20] Martin Reuter, Franz-Erich Wolter, Martha Shenton, and Marc Niethammer. 2009. Laplace–Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer-Aided Design* 41, 10 (2009), 739–755.

[21] Radu Bogdan Rusu, Zoltan Csaba Marton, Nico Blodow, Mihai Dolha, and Michael Beetz. 2008. Towards 3D point cloud based object maps for household environments. *Robotics and Autonomous Systems* 56, 11 (2008), 927–941.

[22] Daniel Scheeres, R Gaskell, S Abe, O Barnouin-Jha, T Hashimoto, J Kawaguchi, T Kubota, J Saito, M Yoshikawa, N Hirata, and others. 2006. The actual dynamical environment about Itokawa. In *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*. 6661.

[23] Philip Shilane and Thomas Funkhouser. 2006. Selecting distinctive 3D shape descriptors for similarity retrieval. In *IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*. IEEE, 18–18.

[24] Sarah Tang and Afzal Godil. 2012. An evaluation of local shape descriptors for 3D shape retrieval. In *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, 82900N–82900N.

[25] Eric Wahl, Ulrich Hillenbrand, and Gerd Hirzinger. 2003. Surflet-pair-relation histograms: a statistical 3D-shape representation for rapid classification. In *3-D Digital Imaging and Modeling, 2003. 3DIM 2003. Proceedings. Fourth International Conference on*. IEEE, 474–481.

[26] Qian Zhang, Jinyuan Jia, and Hongyu Li. 2010. A GPU based 3D object retrieval approach using spatial shape information. In *Multimedia (ISM), 2010 IEEE International Symposium on*. IEEE, 212–219.