COMPUTER-AIDED DESIGN & APPLICATIONS, 2017
VOL. 14, NO. 4, 507–515
http://dx.doi.org/10.1080/16864360.2016.1257192

Computer-Aided Design
Taylor & Francis
Taylor & Francis Group

Learned 3D shape descriptors for classifying 3D point cloud models

Xiaojun Zhao and Horea T. Ilieş

University of Connecticut, USA

ABSTRACT

Recent hardware advances in the field of 3D imaging are democratizing 3D sensing with tremendous scientific and societal implications. Point cloud processing has traditionally required meshing the point clouds output by 3D cameras followed by surface reconstruction, and geometric feature recognition. However, meshing a point cloud is fundamentally an ill-posed problem, and the definition of a good solution is not general. This, in turn, demands new paradigms for processing point cloud information.

In this paper, we focus on the task of classifying 3D point clouds captured with commercial 3D cameras, and we integrate supervised machine learning algorithms with three different yet underexplored shape descriptors, namely Light-Field Descriptors, Angular Radial Transform (ART), and Zernike Descriptors (ZD). We evaluate the classification performance of different machine learning algorithms combined with these different shape descriptors on point cloud models obtained from the Google 3D Warehouse, and we demonstrate good classification performance for 3D point clouds. Furthermore, we show that Zernike Descriptors are practically insensitive to noise levels typically found in point cloud models captured via 3D sensing.

KEYWORDS

Point clouds; classification; 3D shape descriptors; machine learning

1. Introduction

The recent developments in 3D sensing devices that deliver high-quality raw 3D data in real time offer growing opportunities to explore the usage of this data for 3D perception and reasoning tasks. For example, one of the key challenges in developing practical human-robot interaction mechanisms is the automatic interpretation/classification of data collected by the emerging 3D cameras widely used in engineering and science in general, and robotics in particular. This, in turn, demands new paradigms for processing point cloud information, because current point cloud processing algorithms have limited capabilities to automatically extract semantic information from the observed scenes.

Point cloud processing has traditionally required generating a mesh of the point clouds output by the 3D cameras followed by surface reconstruction, and geometric feature recognition. However, meshing a point cloud is fundamentally an ill-posed problem, and the definition of a "good" solution is not general. Furthermore, obtaining good quality meshes from noisy or incomplete point clouds, which could then be used for downstream processing, require frequent user intervention in all practical cases. In turn, this severely limits the range of applications of the 3D sensing and perception technology.

In this paper, we focus on the problem of classifying 3D point clouds captured by commercial 3D cameras, and we are integrating different supervised machine learning classifiers with several capable yet underexplored shape descriptors based on visual similarity (light-field), angular radial transform (ART), and Zernike moments. Specifically, we investigate the use of 3D Zernike descriptors as well as a combination of 2D ART descriptors with light field techniques to construct and compare the performance of practical descriptors for 3D point cloud classification. Importantly, the resulting shape descriptors are invariant under various transformations, as discussed in the section 3 of this paper. We train our classifiers with a database of point clouds corresponding to several common objects obtained by sampling polygonal models obtained from Google's 3D Warehouse and by post-processing them to attain controlled but varying levels of density and noise. We show that these descriptors provide a promising alternative to the current shape descriptors employed for classifying point clouds in the presence of noise. To the best of

CONTACT Xiaojun Zhao xiaojun.zhao@uconn.edu

our knowledge, the work presented in this paper is the first application of 3D Zernike moments to point cloud classification.

The remainder of the paper is structured as follows: section 2 provides an overview of the relevant existing work on existing shape descriptors followed by a detailed explanation on LFD, ART and ZD shape descriptors in Section 3. We report in section 4 the implementation details and evaluate the performance of the implemented shape descriptors on the point cloud database. In section 5 we present our concluding remarks and suggestions for future work.

2. Related work

The classification of shapes relies on the existence of a similarity or dissimilarity measure between shapes. A good representation of the *shape features* in terms of a shape descriptor must be discriminating, efficient to compute and compare, invariant under isometries, insensitive to geometric as well as topologic noise, and robust to degeneracies. In this work we are using the term "shape features" to refer to any collection of attributes of a shape that are distinctive in some sense (such as differential, integral and spectral quantities), which is a rather general definition that captures the largest set of such quantities. Furthermore, we use the term "geometric feature" to refer to the concept of a feature widely accepted in the computer-aided design community, which captures generic shapes and has engineering significance as discussed in [11].

Many 3D shape descriptors have been recently proposed and applied primarily to tessellated models, as detailed in the recent reviews that appear in Tangelder et al. [12] and Kazmi et al. [5]. At the highest level, the descriptors can be grouped based on their representation into: global features (e.g., volume, statistical moments), global feature distributions (e.g., histograms), spatial maps (e.g., spherical harmonics [3]), and local features (e.g., shape spectra) as detailed in [12]. However, almost all existing shape descriptors have been defined for tessellated models, and very few exist that can be applied to native point clouds. For example, Williams et al. [13] developed a practical and convergent estimate of the Laplace-Beltrami operator for point clouds, which is symmetric under real-world conditions, and used it to construct compact shape signatures of point cloud models. These signatures were them used in conjunction with topological clustering techniques via Vietoris-Rips clustering to segment point cloud models of engineering artifacts into geometric features of engineering interest.

Light Field-based techniques have been used in conjunction with 2D signatures to tackle the task of retrieving 3D models from databases (see for example [9]). Furthermore, two-dimensional image moments have been traditionally used for image recognition, but they suffer from noise sensitivity, and information suppression. These difficulties have been addressed in 2D by introducing the Zernike moments defined with Zernike polynomials. For example, Chen et al. [4] proposed the Light Field Descriptors (LFD), which compute 2D Zernike moments and Fourier coefficients based on the silhouettes images taken from cameras on the vertices of a dodecahedron. These Zernike moments have been extended to 3D by Canterakis [2], and have been applied to tessellated model retrieval by Novotni et al. [10], where it is argued that the 3D Zernike moment-based descriptors lead to better retrieval performance and robustness against topological and geometrical artifacts of tessellated models than state of the art descriptors.

Other shape descriptors have been proposed in the literature and applied primarily on polygonal models, but for the sake of brevity we refer the reader to recent surveys such as [5,8,13,15,16].

3. Feature descriptors

In what follows, we present background information about the shape descriptors used in this work, as well as a brief explanation of their key characteristics relevant for classifying point cloud models of rigid objects.

3.1. Angular Radial Transform Descriptors (2D ART)

ART descriptors were proposed by Kim et al. [7]. As a region-based shape descriptor, the original ART is defined as a set of normalized magnitudes of the ART moments or coefficients computed on a 2D image, and is capable of describing both connected and disconnected regions with rotational invariance. These shape descriptors possess several desirable properties, such as compact size, invariance to similarity transformations, and robustness against noise and scaling, and are able to capture features of 2D color images [11]. The ART coefficients, F_{nm} of order n and m, are defined by:

$$F_{nm} = \int_0^{2\pi} \int_0^1 V_{nm}(\rho,\theta) f(\rho,\theta) \rho d\rho d\theta, \qquad (1)$$

where $f(\rho,\theta)$ is the image function expressed in a polar coordinate system, and $V_{nm}(\rho,\theta)$ are the ART basis functions that can be separated in the radial and angular directions:

$$V_{nm}(\rho,\theta) = A_m(\theta) R_n(\rho) \qquad (2)$$

The radial basis function $R_n(\rho)$ and angular basis function $A_m(\theta)$ are defined as:

$$R_n(\rho) = \begin{cases} 1, & n = 0 \\ 2\cos(\pi n\rho), & n \neq 0 \end{cases} \text{ and } A_m(\theta) = \frac{1}{2\pi} e^{jm\theta} \tag{3}$$

Rotational invariance can be obtained by using the magnitude of the ART coefficients [7]. In our experiments, we use $n = 3$ and $m = 12$ to maintain a manageable computational complexity, and ART coefficients are normalized by the magnitude of the ART coefficient F_{00}. In order to achieve translation invariance, the center of the polar coordinate system is defined as the mass center of the object. The ART has been generalized to the indexing of 3D tessellated models (see for example [11]).

3.2. Zernike descriptors (ZD)

Zernike moments are mappings of the function that defines the shape (or image) onto a set of orthogonal polynomials over a unit ball. Similar to the ART descriptor, the Zernike descriptors can be defined from the magnitudes of a set of orthogonal complex moments of objects, and are rotationally invariant.

3.2.1. 2D ZD

2D Zernike polynomials (ZP) in polar coordinates are defined as:

$$V_{\mathrm{nm}}(\rho,\theta) = R_{nm}(\rho)e^{jm\theta}, \tag{4}$$

where $R_{nm}(\rho)$ is the radial polynomial given by [6]:

$$R_{\mathrm{nm}}(\rho) = \sum_{s=0}^{(n-|m|)/2} (-1)^s \times \frac{(n-s)!}{s!\left(\frac{n+|m|}{2}-s\right)!\left(\frac{n-|m|}{2}-s\right)!}\rho^{n-2s} \tag{5}$$

In the above two equations, order n is a positive integer or 0, and repetition m is an integer such that $|m| \leq n$ and $n - m$ is even. 2D Zernike moments can be constructed in the following discrete form:

$$A_{\mathrm{nm}} = \frac{n+1}{\pi} \sum_x \sum_y f(x,y) V_{nm}{}^*(\rho,\theta), \quad \text{for } x^2 + y^2 \leq 1 \tag{6}$$

Similar to ART descriptors, $|A_{\mathrm{nm}}|$ are rotation invariant and can be used as shape descriptors.

3.2.2. 3D ZD

Following the notation used in [10], 3D Zernike polynomials Z_{nl}^m in polar coordinates are given by:

$$Z_{nl}^m(r,\theta,\phi) = R_{nl}(r) \cdot Y_l^m(\theta,\phi) \tag{7}$$

and restricting l such that $|l| \leq n$ and $n - l$ be an even number. In equation (7) $R_{nl}(r)$ are the radial polynomials and Y_l^m are spherical harmonics. In the Cartesian coordinate system, the 3D Zernike polynomials can be represented as:

$$Z_{nl}^m(x,y,z) = \sum_{v=0}^{k} q_{kl}^v r^{2v} e_l^m(x,y,\mathrm{z}), \tag{8}$$

where q_{kl}^v are coefficients, $2k = n - l$ and the harmonic polynomials e_l^m are defined as $e_l^m = r^l Y_l^m(\vartheta,\varphi)$ in spherical coordinates. These harmonic polynomials can also be expressed in Cartesian coordinates as:

$$e_l^m(x,y,z) = \frac{\sqrt{(2l+1)(l+m)!(l-m)!}}{l!} r^l \left(\frac{ix-y}{2}\right)^m \times z^{l-m} \cdot \sum_{\mu=0}^{\left\lfloor \frac{l-m}{2} \right\rfloor} \binom{l}{\mu}\binom{l-\mu}{m+\mu} \times \left(-\frac{x^2+y^2}{4z^2}\right)^\mu \tag{9}$$

The 3D Zernike moments Ω_{nl}^m of an object function f are defined as [10]:

$$\Omega_{nl}^m := \frac{3}{4\pi} \int_{|x|\leq 1} f(\boldsymbol{x}) \overline{Z_{nl}^m(\boldsymbol{x})} d\boldsymbol{x} \tag{10}$$

A rotationally invariant version of these moments can be obtained by normalizing the moment vector $\Omega_{nl} = (\Omega_{nl}^l, \Omega_{nl}^{l-1}, \ldots, \Omega_{nl}^{-l})$. Later in the paper we explore the impact of the orders of the 3D Zernike descriptors on the performance of the classification as well as its robustness against different level of noise.

3.3. Light field descriptors (LFD)

LFD uses the observation that two similar objects look similar from similar viewing angles. We set 20 viewpoints (or cameras) on 20 vertices of a regular dodecahedron. Since the cameras on the opposite vertices would produce the same silhouettes, 10 object views are needed for each model. To achieve the rotational invariance property, each 3D model is "observed" by 10 cameras in 10 different orientations. Therefore, a total of 100 silhouettes are determined for each model, and each 2D silhouette is encoded by a feature vector extracted by

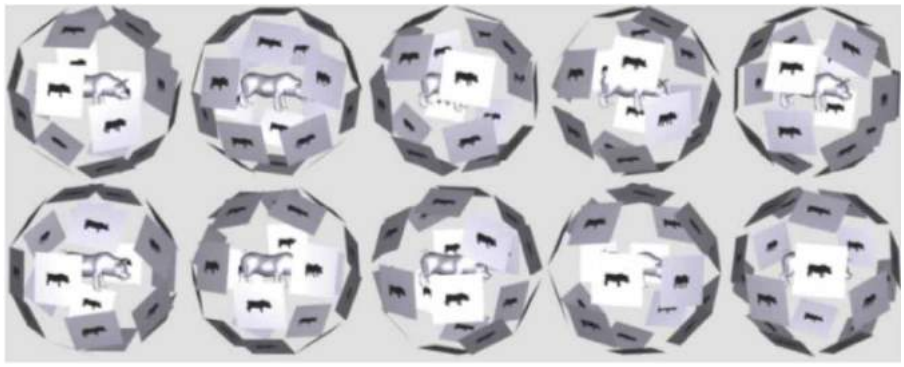

Figure 1. A set of LFDs for a 3D model.

ART. Fig. 1, reproduced from [4], illustrates the set of LFDs used for one model to obtain rotational invariance.

Consequently, a total of 100 silhouettes are rendered for each model, and each image is encoded by a feature vector with 35 2D Zernike moments or ART moments. Therefore, the length of LFD+2D ART/2D ZD is 3500. We use PCA to obtain a more compact descriptor for each model.

4. Experimentations and results

4.1. Experimental setup

We create our own 3D model database as the test bed for the experimentation by using models downloaded from Google's 3D Warehouse. However, we note that there are other object databases that have been used primarily for processing meshed models, such as the Princeton, COSEG and NTU databases. Our database currently contains 185 models belonging to 6 common object categories, including 30 cars, 29 planes, 31 mugs, 31 tables, 34 chairs and 30 desk lamps. These categories of objects ensure a geometrically diverse database. A sample of the 3D meshed models that we used to build our dataset is shown in Fig. 2: each row represents a model class, and 5 models shown in each row are selected arbitrarily from the corresponding category.

The pipeline of our experimentations is shown in Fig. 3. It consists of the following three modules:

1. Model Preprocessing:
   first step is to convert the meshed models into point clouds models via Poisson disk sampling, and to estimate the surface normal at each point from the k-nearest neighbors of each query point. Then, the model is scaled into a unit cube and then translated to the origin of the coordinate system. Next object views are created for point cloud models to compute LFD, and point cloud models are voxelized for 3D ZD.

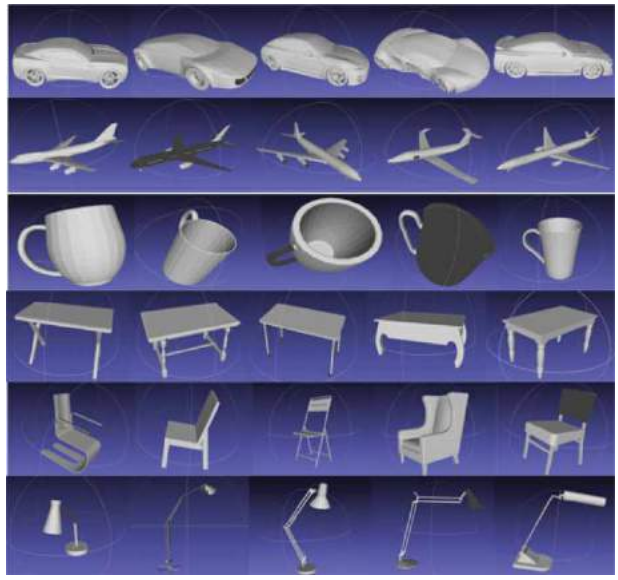

Figure 2. A sample of the 3D meshed models downloaded from Google's 3D Warehouse.

2. Feature Extraction:
   feature vector for each model is computed by using LFD (with 2D ZD and ART) and 3D ZD. This allows the construction of a feature matrix having a size of $185 \times L$, where L is the length of a single feature vector. To reduce the high computational complexity of training machine learning classifiers, we use PCA to reduce the dimensionality of the feature vectors to 20-30 feature components.
3. Training and Testing:
   partition the dataset into training set and test set. The training set contains 125 models, the remaining 60 are used as a test set. To prevent the overfitting issue, we repeat the whole process 50 times (with and without the k-fold cross validation) with different partitions of the data set. The final step is to compare the classification effectiveness and performance among different classifiers.

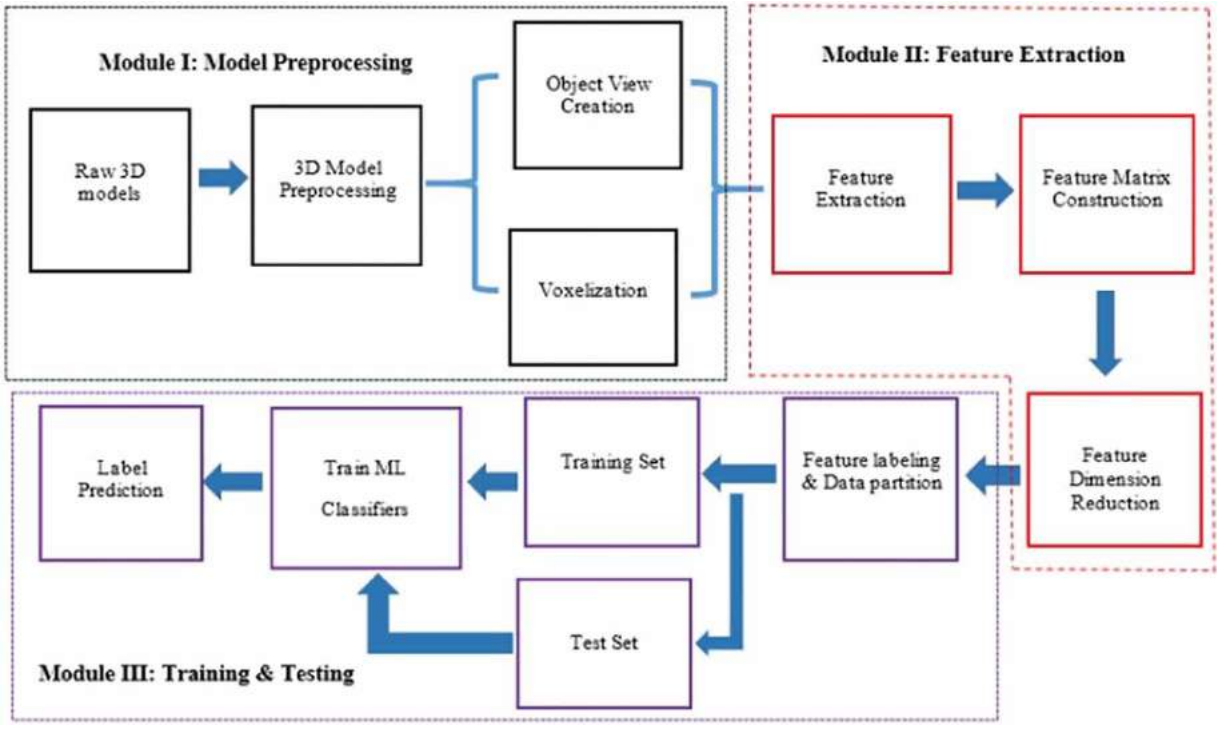

Figure 3. Pipeline of 3D model classification system.

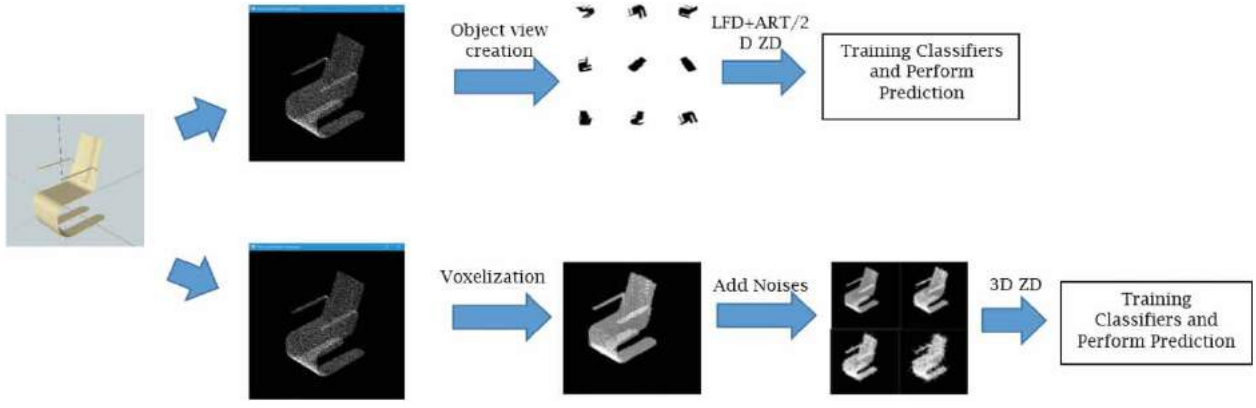

Figure 4. Illustration of the classification procedure for a chair model.

A specific example of the above workflow for a chair model is illustrated in the Fig. 4.

4.2. Experimental results

We implemented the LFD+ART/2D ZD and 3D ZDs in our experiments. To evaluate the classification performance, we first compare the overall classification accuracies among different implemented classifiers with and without cross validation, and then compare the performance for each model category. The classification process was repeated 50 times, each time with randomly partitioning the data into training and test sets. The average classification reported below is the average classification for all 50 runs. In addition, we also investigate how the 3D ZDs respond to point cloud noise that is similar with that seen by commercial 3D sensing devices.

4.2.1. LFD implementation and results

Fig. 5 shows a number of object views for several models selected form our dataset.

After the object views are created, we use 2D ZD and 2D ART descriptors to extract features of the rendered silhouettes of the point clouds. For 2D ZD, we compute the first 35 moments (up to 7^{th} order), while obtaining the same number of ART moments by using 3 radial functions and 12 angular functions ($n < 3, m < 12$). The comparisons of average classification accuracy for LFD+2D

Figure 5. Object views of some selected models.

Table 1. Classifcation accuracy by using LFD+2D ZD and LFD+2D ART with three different machine learning algorithms: Multi-Layer Perceptron (MLP), k-Nearest Neighbor (kNN) and Random Forests.

		Avg. Classification Accuracy	
Light-Field Descriptors		With 10-fold cv (repeated 50 times)	Without cv (repeated 50 times)
MLP (Input layer size: 40; Hidden layer size: 25; Training Function: Sigmoid)	LFD+2D ZD	81.62%±1.0%	79.33%±1.0%
	LFD+ART	80.54%±1.0%	80.00%±1.0%
KNN (Feature size: 30; k = 1)	LFD+2D ZD	98.19%±1.0%	97.70%±1.0%
	LFD+ART	98.22%±1.0%	98.57%±1.0%
Random Forests (#Trees: 120)	LFD+2D ZD	98.16%±1.0%	97.00%±1.0%
	LFD+ART	98.70%±1.0%	98.08%±1.0%

ZD and LFD+2D ART among different classifiers are presented in Tab. 1.

We observe that LFD+ART slightly outperform LFD+2D ZD overall, and that the kNN and Random Forests perform better than the MLP among the implemented classifiers.

Our experiments with the object database described above show that the classification accuracy is relatively uniform across the model categories as shown in Table 2, where we provide the classsification accuracy of the kNN classifier.

4.2.2. 3D ZD Implementation and Results

In order to construct the 3D Zernike descriptors directly on the point cloud models, we first convert each meshed 3D models to a point cloud using Poisson-disk sampling, then voxelize the point cloud and select those voxels that contain a single point of the point cloud. We note that the result of this voxelization is sensitive to noise in the point cloud, but it is computationally efficient.

To investigate the impact of the order of the Zernike polynomial on the classification performance, we select n to be 15, 20, 25 and 30. Each polynomial order corresponds to different number of feature components. The results are summarized in the Tab. 3.

Table 3 shows that the increase of the polynomial order does not necessarily improve the classification accuracy. Clearly, the higher the polynomial order, the longer it takes to compute the features of a 3D point cloud model. We observeed that a good compromise between the classification performance and the associated computational cost for models in our database is achieved for polynomial orders between 20 and 25.

The classification quality for each class can be plotted in the Receiver Operating Characteristic (ROC) curve (true positives vs false positives rate). The more the curve is leaning toward the upper left corner (larger area under curve), the better the result is.

Fig. 6 shows the ROC curve of the Random Forest Classifier. We can see that the area under the ROC curve for each object category is in the range of

Table 2. Classification results for each category by using the nearest neighbor classifier.

	LFD+ZD			LFD+ART		
Light Field Descriptors	Avg. no. of Correct Classified Models	Avg. no. of models in the test set	Avg. Accuracy	Avg. no. of Correct Classified Models	Avg. no. of models in the test set	Avg. Accuracy
Chair	10.84	11.32	95.96%	10.78	11.14	96.76%
Table	9.88	10.30	95.92%	10.22	10.48	97.50%
Mug	10.34	10.34	100.00%	10.38	10.38	100.00%
Plane	9.36	9.36	100.00%	9.30	9.30	100.00%
Car	9.36	9.36	100.00%	9.30	9.30	100.00%
Lamp	8.84	9.32	94.84%	9.16	9.40	97.44%
Total	58.62	60.00	97.70%	59.14	60.00	98.57%

Table 3. Performance comparisons among using 4 different orders of ZP.

3D ZD	*n = 15 (72 feature components)*	*n = 20 (121 feature components)*	*n = 25 (182 feature components)*	*n = 30 (256 feature components)*
KNN (k = 1)	87.87%	86.43%	86.93%	87.63%
Random Forest (# tree = 120)	82.37%	83.06%	85.17%	82.90%
MLP	75.00% hidden layer size:10	73.33% hidden layer size:10	71.67% hidden layer size:40	70.00% hidden layer size:40
RBF Network	85.67%	88.17%	89.17%	89.33%
Adaptive Boosting	76.17%	74.75%	75.75%	73.00%
kNN (5-fold cv)	89.52%	89.03%	89.14%	90.48%
RBF network (5-fold cv)	89.18%	92.49%	93.63%	92.97%

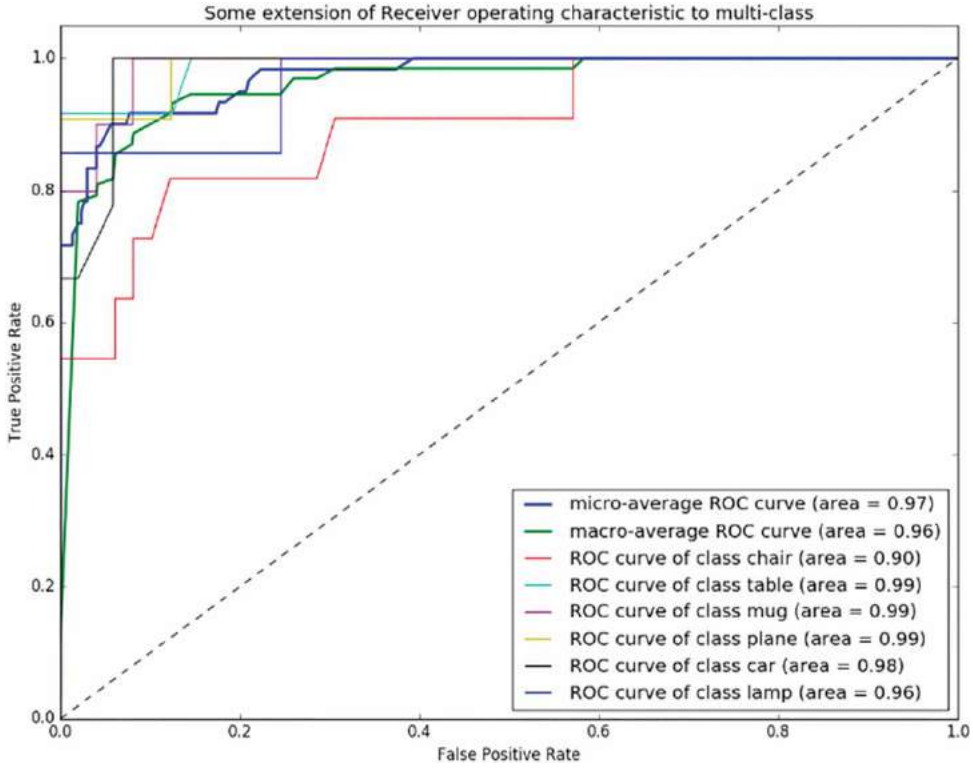

Figure 6. ROC curve of each class by using Random Forest Classifier.

0.90 to 0.99, which demonstrates a good performance of 3D ZDs.

In real world scenarios, the models obtained by 3D imaging devices inevitably contain noise. Therefore, a desirable shape descriptor should be insensitive to noise in a certain acceptable range. To investigate the ZD's sensitivity to noise, we created four levels of point cloud noises (including level 0, which denotes a noise-less model) by randomly selecting one-third of the total points in each model and adding random noise along the normal direction for each such point. Each level can be represented by a noise coefficient C that is proportional to the noise levels. Figure 7 shows a point cloud model with different levels of noise.

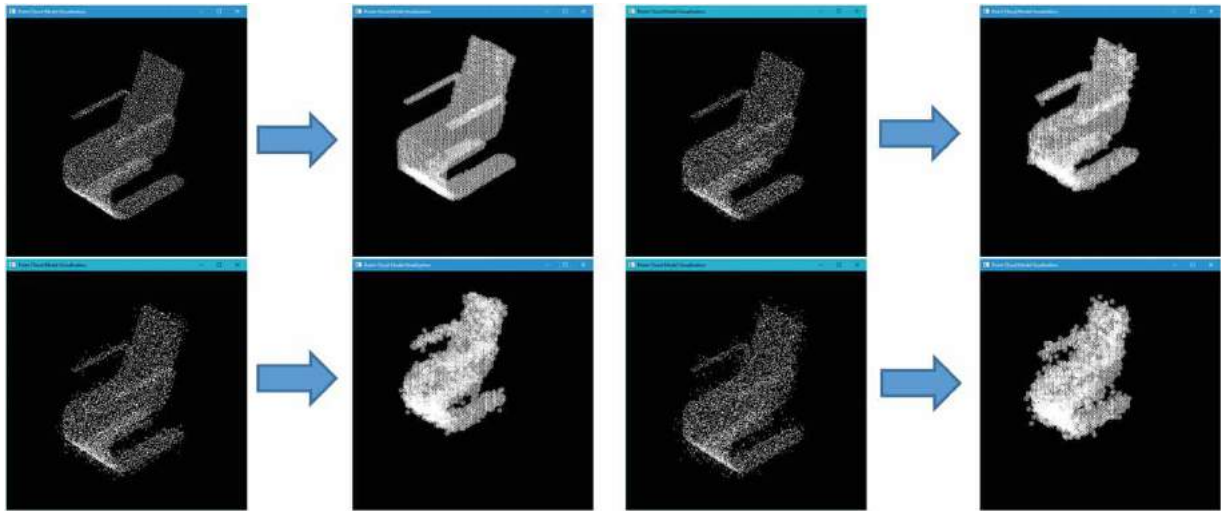

Figure 7. Addition of noise to the model (from top left to right, bottom left to right): (a) noise level 0 with C = 0.00, (b) noise level 1 with C = 0.03, (c) noise level 2 with C = 0.06, (d) noise level 3 with C = 0.09.

Table 4. Classification performance for different levels of noise (with $n = 25$).

Noise Level	Level 0	Level 1	Level 2	Level 3
MLP	71.67%	70.00%	73.33%	70.00%
RBF Network	88.67%	87.83%	89.67%	88.50%
KNN (K = 1)	88.13%	87.17%	88.50%	86.93%
Random Forest	85.17%	83.07%	83.50%	83.83%

We use 3D ZD with an order of 25 to determine the classification performance for models that have these four different levels of noise, and the results are shown in the Tab 4.

The data shown in Table 4 indicates that the classification accuracy does not change significantly as the level of noise increases. This suggests that the 3D Zernike descriptors are robust against random noise levels similar to those observed in 3D sensing with commercial cameras.

5. Conclusions

This work explores the classification task for 3D point cloud models by incorporating supervised machine learning approaches with powerful shape descriptors that have traditionally been used for classifying polygonal models. We consider two different approaches to feature extraction from point clouds, namely Light-Field Descriptors built with either 2D Zernike or 2D ART moments, and 3D Zernike Descriptors computed directly on the point clouds. The major difference between the two types of approaches is that the LFD-based approaches rely on view-similarities and extract lower-dimensional features (i.e. 2D shapes), while 3D ZD compute features directly based on 3D data (i.e. 3D point clouds). Our preliminary experimental results showed that LFD+ART or LFD+2D ZD outperform 3D ZD in terms of classification performance, but they all have the potential to robustly and effectively classify 3D point cloud models without requiring a mesh of the point cloud. Furthermore, our experiments show that 3D Zernike descriptors are robust against noise levels typically found in point cloud data output by current commercial RGB-D cameras.

Our preliminary experiments show that the LFD+2D ZD/ART and the 3D Zernike descriptors provide a promising alternative to the current shape descriptors employed for classifying noisy point clouds. Furthermore, the practicality of 3D Zernike descriptors coupled with their potential for parallel implementations on the GPU [1] makes them capable candidates for real-time applications in 3D sensing and perception.

Acknowledgements

This work was supported in part by the National Science Foundation grants CMMI-1200089, CMMI 1462759, and CHS 1526249.

ORCID

Xiaojun Zhao http://orcid.org/0000-0002-7769-5272

References

[1] Berjón, D.; Arnaldo, S.; and Morán, F.: A parallel implementation of 3D Zernike moment analysis, IS&T/SPIE Electronic Imaging, International Society for Optics and Photonics, 7872, 2011. http://dx.doi.org/10.1117/12.876683

[2] Canterakis, N.: 3D Zernike moments and Zernike affine invariants for 3D image analysis and recognition, 11th Scandinavian Conf. on Image Analysis. 1999.
[3] Barra, V.; Biasotti, S.: 3D shape retrieval using kernels on extended Reeb graphs, Pattern Recognition, 46(11), 2013, 2985–2999. http://dx.doi.org/10.1016/j.patcog.2013.03.019
[4] Chen, D. Y.; Tian, X. P.; Shen, Y. T.; and Ouhyoung, M.: On visual similarity based 3D model retrieval, Computer Graphics Forum, 22(3), 2003, 223–232. http://dx.doi.org/10.1111/1467-8659.00669
[5] Kazmi, I.K.; Lihua, Y.; Zhang, J.J.: A survey of 2D and 3D shape descriptors. IEEE 10th International Conference on Computer Graphics, Imaging and Visualization (CGIV), 2013, 1–10. http://dx.doi.org/10.1109/cgiv.2013.11
[6] Khotanzad, A.; Yaw, H. H.: Invariant image recognition by Zernike moments, IEEE Transactions on Pattern Analysis and Machine Intelligence, 12(5), 1990, 489–497. http://dx.doi.org/10.1109/34.55109
[7] Kim, W.-Y.; Kim, Y.-S.: A new region-based shape descriptor, ISO/IEC MPEG99/M5472 Maui, Hawaii, 1999.
[8] López, G.L.; Negrón, A.P.P.: Jiménez, A.D.A.; Rodríguez, J.R.; Paredes, R.I.: Comparative analysis of shape descriptors for 3D objects, Multimedia Tools and Applications, 2016, 1–48. http://dx.doi.org/10.1007/s11042-016-3330-5
[9] Napoléon, T.; Sahbi, H.: From 2D silhouettes to 3D object retrieval: contributions and benchmarking, Journal on Image and Video Processing, 2010(1), 2010. http://dx.doi.org/10.1155/2010/367181
[10] Novotni, M.; Klein, R.: 3D Zernike descriptors for content based shape retrieval, Proceedings of the eighth ACM symposium on Solid Modeling and Applications, ACM, 2003, 216–225. http://dx.doi.org/10.1145/781606.781639
[11] Ricard, J.; Coeurjolly, D.; Baskurt, A.: Generalizations of angular radial transform for 2D and 3D shape retrieval, Pattern Recognition Letters, 26(14), 2005, 2174–2186. http://dx.doi.org/10.1016/j.patrec.2005.03.030.
[12] Shah, J.J.: Conceptual development of form features and feature modelers, Research in Engineering Design, 2(2), 1991, 93–108. http://dx.doi.org/10.1007/BF01579254
[13] Tangelder, J.W.H.; Veltkamp, R.: A survey of content based 3D shape retrieval methods, Multimedia tools and applications, 39(3), 2008, 441–471. http://dx.doi.org/10.1007/s11042-007-0181-0
[14] Williams, R.M.; Ilies, H.T.: Practical Shape Analysis and Segmentation Methods for Point Cloud Models. Technical report, preprint, 2016.
[15] Yang, M.; Kpalma, K.; Ronsin, J.: A survey of shape feature extraction techniques, Pattern recognition, 2008, 43–90. http://dx.doi.org/10.1007/10.5772/6237
[16] Zhang, L.; da Fonseca, M. J.; Ferreira, A.: Survey on 3D shape descriptors, DecorAR, Technical Report, POSC/EIA/59938/2004, 2007.