

Fine-Grained 3D Shape Classification With Hierarchical Part-View Attention

Xinhai Liu, Zhizhong Han^{id}, Yu-Shen Liu^{id}, *Member, IEEE*, and Matthias Zwicker^{id}, *Associate Member, IEEE*

Abstract—Fine-grained 3D shape classification is important for shape understanding and analysis, which poses a challenging research problem. However, the studies on the fine-grained 3D shape classification have rarely been explored, due to the lack of fine-grained 3D shape benchmarks. To address this issue, we first introduce a new 3D shape dataset (named FG3D dataset) with fine-grained class labels, which consists of three categories including airplane, car and chair. Each category consists of several subcategories at a fine-grained level. According to our experiments under this fine-grained dataset, we find that state-of-the-art methods are significantly limited by the small variance among subcategories in the same category. To resolve this problem, we further propose a novel fine-grained 3D shape classification method named FG3D-Net to capture the fine-grained local details of 3D shapes from multiple rendered views. Specifically, we first train a Region Proposal Network (RPN) to detect the generally semantic parts inside multiple views under the benchmark of generally semantic part detection. Then, we design a *hierarchical part-view attention aggregation* module to learn a global shape representation by aggregating generally semantic part features, which preserves the local details of 3D shapes. The part-view attention module hierarchically leverages part-level and view-level attention to increase the discriminability of our features. The part-level attention highlights the important parts in each view while the view-level attention highlights the discriminative views among all the views of the same object. In addition, we integrate a Recurrent Neural Network (RNN) to capture the spatial relationships among sequential views from different viewpoints. Our results under the fine-grained 3D shape dataset show that our method outperforms other state-of-the-art methods. The FG3D dataset is available at <https://github.com/liuxinhai/FG3D-Net>.

Index Terms—Fine-grained shape classification, 3D objects, generally semantic part, dataset, attention, recurrent neural network.

I. INTRODUCTION

LEARNING a shape representation from multiple rendered views is an effective way to understand 3D shapes [1]–

Manuscript received February 12, 2020; revised August 21, 2020 and December 5, 2020; accepted December 22, 2020. Date of publication January 8, 2021; date of current version January 14, 2021. This work was supported in part by the National Key Research and Development Program of China under Grant 2020YFF0304100 and Grant 2018YFB0505400, in part by the National Natural Science Foundation of China under Grant 62072268, in part by the Tsinghua-Kuaishou Institute of Future Media Data, and in part by NSF under Award 1813583. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Lucio Marcenaro. (*Corresponding author: Yu-Shen Liu.*)

Xinhai Liu and Yu-Shen Liu are with the School of Software, BNRist, Tsinghua University, Beijing 100084, China (e-mail: lxh17@mails.tsinghua.edu.cn; liuyushen@tsinghua.edu.cn).

Zhizhong Han and Matthias Zwicker are with the Department of Computer Science, University of Maryland, College Park, MD 20737 USA (e-mail: h312h@mail.nwpu.edu.cn; zwicker@cs.umd.edu).

Digital Object Identifier 10.1109/TIP.2020.3048623

1941-0042 © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.

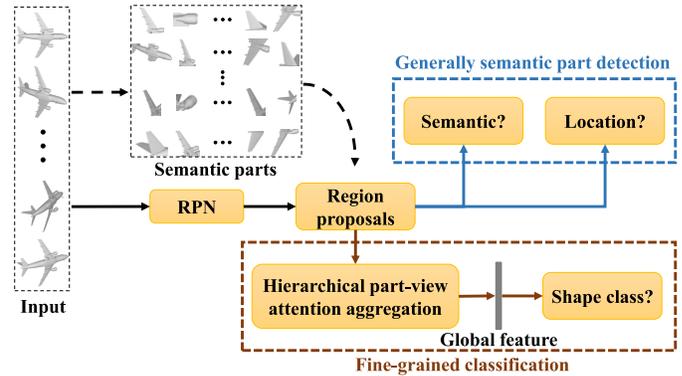


Fig. 1. The illustration of our FG3D-Net. In FG3D-Net, we first render each 3D shape into multiple views that are propagated into a Region Proposal Network (RPN) [6] to generate a set of region proposals. Then, the extracted region proposals are used in two different branches including generally semantic part detection (top branch in blue box) and fine-grained classification (bottom branch in brown). In the generally semantic part detection branch, we predict the semantic score and the bounding box location for each region proposal. According to the semantic scores, several region proposals are selected to extract the global feature of the input 3D shape in the fine-grained classification branch. Specifically, we introduce a novel module named *hierarchical part-view attention aggregation* to effectively capture the fine-grained 3D shape details for fine-grained classification.

[4]. Influenced by the great success of Convolutional Neural Networks (CNNs) in the recognition of 2D images under large-scale datasets, such as ImageNet [5], 2D CNNs are intuitively applied to learn the representation for 3D shapes. For example, the pioneering MVCNN [1] first projects a 3D shape into multiple views from different viewpoints and then obtains a global 3D shape representation by aggregating the view features with a view pooling layer, where the view features are extracted by a shared CNN. Previous view-based methods have achieved satisfactory performance for 3D shape recognition under large variance among different categories. However, it is still nontrivial for these methods to capture the small variance among subcategories in the same category, which limits the discriminability of learned features for fine-grained 3D shape recognition.

Fine-grained 3D shape recognition, which aims to discriminate 3D shapes within the same category, such as airliners, fighters and seaplanes within the airplane category, is quite challenging. Specifically, there are two major issues that limit the performance of fine-grained classification of 3D shapes. On the one hand, many large-scale image datasets (e.g. CUB-200-2011 [7] and Stanford dog dataset [8]) have been developed for fine-grained object classification and recognition in 2D area, but large-scale 3D object datasets are still eagerly needed for

fine-grained 3D shape classification. Recently, several well-known large-scale 3D shape repositories, such as ShapeNet [9] and ModelNet [10], have been developed for learning shape representations in various applications. While they contain large numbers of 3D objects from different categories, they are not well organized for fine-grained 3D shape classification. On the other hand, existing view-based methods for learning 3D shape representations are still suffering from capturing the fine-grained details from multiple views in a more comprehensive way, which is necessary for fine-grained classification of 3D shapes. Intuitively, subtle and local differences are usually exposed in parts on the objects, so it is vital to leverage the parts in fine-grained 3D shape classification.

Previous multi-view methods such as MVCNN [1] and RotationNet [3] usually extract a feature from pixel-level information in each view first, and then aggregate the extracted view features into a global shape representation. However, these methods do not capture part-level semantics from multiple views. To address this problem, our recent work named Parts4Feature [11] utilized a region proposal network to detect generally semantic parts from multiple views and then learned to directly aggregate all generally semantic parts into a global shape representation. However, there are still several unresolved issues in Parts4Feature, which limits its performance in fine-grained 3D shape classification. First, Parts4Feature cannot capture the correlation of generally semantic parts in the same view, which makes it unable to filter out the meaningless generally semantic parts. Second, Parts4Feature ignores view-level information such as the importance of each view and the spatial relationship among sequential views, which is important for learning the fine-grained 3D shape features.

To solve the above-mentioned issues, in this paper we propose a novel fine-grained 3D shape classification method, named FG3D-Net, as shown in Fig. 1, which leverages a hierarchical part-view attention aggregation module to capture the fine-grained features. Similar to [11], we first employ a region proposal neural network to detect generally semantic parts in each one of multiple views, which is considered to contain rich local details of 3D shapes. By introducing the supervision information of bounding boxes from other 3D segmentation datasets, our FG3D-Net is able to explore the fine-grained details inside local parts. Then, to aggregate all these extracted generally semantic parts, we leverage semantic information at different levels including part-level, view-level, and shape-level. Specifically, we introduce part-level attention to highlight the important parts in each view and view-level attention to highlight discriminative views among all the views of the same object. To take advantage of sequential input views as used in [12], we employ a Recurrent Neural Network (RNN) to encode the spatial relationship among views. In order to eliminate the impact of the initial view in the RNN inputs, we integrate a global 3D shape feature with a max-pooling operation as [1], which is invariant to the permutation of views.

In addition, we introduce a new fine-grained 3D shape dataset that consists of three object categories including *Airplane*, *Car*, and *Chair*, where dozens of subcategories are

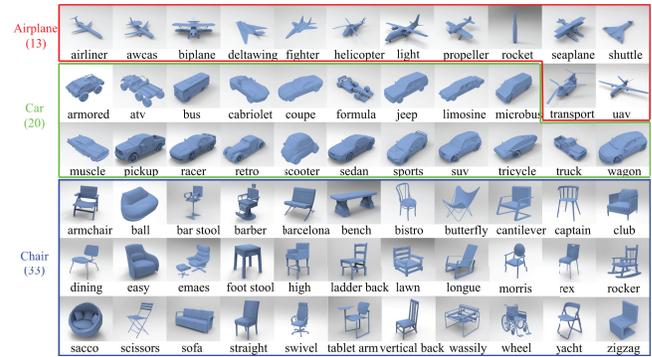


Fig. 2. There are three shape categories in our fine-grained dataset including *Airplane*, *Car* and *Chair*. Specifically, 13 shape subcategories are included in the *Airplane* category such as airliner, fighter and seaplane, 20 shape subcategories such as bus, jeep and scooter, are involved in the *Car* category, and the *Chair* category consists of 33 shape subcategories including bistro, captain, rocker, etc.

constructed in each category. All 3D objects in the dataset are collected from several online repositories and are organized under the WordNet [13] taxonomy. Different from existing 3D datasets such as ShapeNet [9] and ModelNet [10], our dataset is organized as a fine-grained 3D shape classification benchmark, where each 3D object is strictly assigned to one single subcategory in its category. Our main contributions are summarized as follows.

- We present a new fine-grained 3D shape dataset (named FG3D dataset) consisting of three categories including *Airplane*, *Car* and *Chair*, which contains tens of thousands of 3D shapes with unique sub-category labels. This enables the learning of fine-grained features for fine-grained 3D shape classification.
- We propose a novel deep neural network named FG3D-Net to extract a global 3D shape representation that captures the fine-grained local details from generally semantic parts. FG3D-Net further includes a part-level and view-level attention mechanism to highlight the more semantic generally semantic parts in each view and the more distinctive views for each object, respectively.
- We show that FG3D-Net outperforms state-of-the-art techniques in the fine-grained 3D shape classification task.

II. RELATED WORK

A. 3D Shape Datasets

3D shapes are widely used in various applications, such as robotics [14] and 3D modeling [15]. In recent years, 3D shape understanding has attracted a lot of research interest. However, due to the inherent complexity of 3D shapes, 3D shape understanding is still a challenging problem in 3D computer vision. Benefiting from deep learning models, deep neural network based methods have achieved significant performance in 3D shape recognition. These methods require large-scale 3D datasets which are crucial for training deep neural networks and evaluating their performance. Researchers have been working on building some large-scale repositories

[10], [16]–[18], which are widely adopted to evaluate deep neural networks in various applications. With the development of web-scale datasets, ShapeNet [9] has collected a large-scale set of synthetic 3D CAD models from online open-sourced 3D repositories, including more than three million models and three thousand object categories. Some other 3D repositories [19]–[22] were also proposed, which contain semantic labels for the segmented components of 3D shapes. Recently, PartNet [23] provided more fine-grained part annotations to support fine-grained 3D shape segmentation tasks. However, there is still no suitable 3D shape benchmark for the fine-grained 3D shape classification task so far.

To address this problem, some previous studies introduced several fine-grained image datasets and evaluated the results of fine-grained 3D shape classification. FGVC-Aircraft [24] has collected ten thousand images of aircraft spanning 100 aircraft objects. The Car dataset [25] contains 16,185 images from 196 subcategories of cars. Unfortunately, previous 3D fine-grained datasets usually represent each 3D shape with a single 2D image, which can be regarded as the benchmarks of fine-grained 2D image classification. Therefore, in this paper, we present a fine-grained 3D shape dataset containing 3D shapes represented by 3D meshes, where these shapes can be easily translated into other 3D data formats such as rendered views, point clouds and volumetric voxels. In Fig. 2, we show all the subcategories in our fine-grained dataset of three categories including airplanes, cars and chairs, respectively. The construction of the fine-grained dataset will be illustrated in Section III. FG3D DATASET. With the help of this proposed fine-grained 3D shape dataset, we can evaluate the ability of algorithms in capturing shape details under fine-grained 3D shape classification.

B. Fine-Grained Classification

Fine-grained classification aims to classify many subcategories under a same basic-level category such as different cats, dogs or cars. Due to the large intra-subcategory variance and the small inter-subcategory variance, fine-grained classification is a long standing problem in computer vision. Recently, deep learning based methods have been widely applied to fine-grained image classification and achieved significant improvement over traditional methods. From the perspective of fine-grained image classification, current methods can be summarized into three categories: (1) ensemble of networks based methods, (2) visual attention based methods, (3) part detection based methods.

Firstly, ensemble of networks based methods were proposed to learn different representations of images for better classification performance with multiple neural networks. MGD [26] trained a series of CNNs at multiple levels, which focus on different regions of interests in images. B-CNN [27] was proposed with a bilinear CNN model, which jointly combined two CNN feature extractors. Spatial Transformers [28] were proposed with a learnable model that consists of three parts including localization network, grid generator, and sampler. The discriminative parts inside images were captured by four parallel spatial transformers on images and passed to the part description subnets.

Secondly, influenced by attention mechanisms, researchers also focus on searching discriminative parts dynamically, rather than dealing with images directly. AFGC [29] employed an attention mechanism for fine-grained classification system, which utilized the information of multi-resolution corps to obtain the location and the object on the input images.

Finally, subtle and local differences are usually shown in discriminative parts of objects. Therefore, discriminative part detection is very important for fine-grained shape classification. The R-CNN family approaches [6] employed a popular strategy that first generates thousands of candidate proposals and then filters out these proposals with confidence scores and bounding box locations. Zhang *et al.* [30] proposed to detect discriminative parts for fine-grained image classification and trained a classifier on the features of detected parts. Recently, some studies [31], [32] focused on detecting discriminative parts under the weakly supervised setting, which means neither object nor part annotations are needed in both training and testing phases. In [31], part detectors were trained by finding constellations of neural activation patterns computed using convolutional neural networks. Specifically, the neural activation maps were computed as part detectors by using the outputs of a middle layer of CNN. All these methods have been proposed to accomplish fine-grained classification of 2D images. However, the fine-grained classification of 3D shapes has been rarely explored so far.

To address this issue, we propose FG3D-Net to learn fine-grained global 3D shape features by capturing geometry details in generally semantic parts.

C. Deep Learning Based Methods for 3D Shapes

Benefiting from the advances in deep learning, deep learning based methods have achieved significant performance in 3D shape understanding tasks such as shape classification and recognition. In general, current methods can be categorized into mesh-based, voxel-based, point cloud-based, and view-based deep learning methods. To directly learn 3D features from 3D meshes, circle convolution [33] and mesh convolution [34] were proposed to learn local or global features. Similar to images, voxels also have a regular structure that can be learned by deep learning models, such as CRBM [35], SeqXY2SeqZ [36] and DSO-Net [37], fully convolutional denoising autoencoders [38], CNNs [39] and GANs [40]. These methods usually employ 3D convolution to better capture the contextual information inside local regions. Moreover, Tags2Parts [41] discovered semantic regions that strongly correlate with user-prescribed tags by learning from voxels using a novel U-Net. As a series of pioneering work, PointNet [42] and PointNet++ [43] inspired various supervised methods [44]–[48] to understand point clouds. Through self-reconstruction, FoldingNet [49] and LatentGAN [50]–[52] learned global features with different unsupervised strategies. Similar to the light field descriptor (LFD), GIFT [2] measured the difference between two 3D shapes using their corresponding view feature sets. Moreover, pooling panorama views [53], [54] or rendered views [1], [4], [11] are more widely used to learn global features. Different improvements from camera trajectories [55],

TABLE I
THE STATISTICS OF OUR FG3D DATASET WHICH CONSISTS OF 3 CATEGORIES AND 66 SUBCATEGORIES

Subcategory	Train	Test	Total	Subcategory	Train	Test	Total	Subcategory	Train	Test	Total
airliner	955	100	1055	muscle	468	100	568	dining	614	100	714
awcas	24	10	34	pickup	201	50	251	easy	990	100	1090
biplane	127	50	177	racer	383	100	483	emaes	73	40	113
deltawing	186	50	236	retro	102	50	152	foot stool	186	50	236
fighter	545	100	645	scooter	25	5	30	high	256	100	356
helicopter	549	100	649	sedan	2027	100	2127	ladder back	241	50	291
light	101	30	131	sports	427	100	527	lawn	84	30	114
propeller	288	100	388	suv	209	100	309	longue	665	100	765
rocket	390	100	490	tricycle	14	5	19	morris	154	50	204
seaplane	35	20	55	truck	131	50	181	rex	407	100	507
shuttle	201	50	251	wagon	427	100	527	rocker	121	50	171
transport	15	7	22	Car total	7010	1315	8325	sacco	114	50	164
uav	25	15	40	armchair	1578	100	1678	scissors	65	20	85
Airplane total	3441	732	4173	ball	226	100	326	sofa	307	100	407
armored	16	5	21	bar	606	100	706	straight	1553	100	1653
atv	77	30	107	barber	20	5	25	swivel	777	100	877
bus	936	100	1036	barcelona	35	10	45	tablet armed	51	20	71
cabriolet	477	100	577	bench	41	20	61	vertical back	748	100	848
coupe	528	100	628	bistro	48	20	68	wassily	18	5	23
formula	62	20	82	butterfly	10	5	15	wheel	7	5	12
jeep	276	100	376	cantilever	310	100	410	yacht	55	20	75
limousine	109	50	159	captain	217	50	267	zigzag	89	30	119
microbus	115	50	165	club	458	100	558	Chair total	11124	1930	13054

TABLE II
THE DATASET COMPARISON BETWEEN OUR FG3D AND SHAPENET

Category	#Subcategories	#Overlap Subcategories	#Total Shapes	#Overlap Shapes
Airplane (FG3D)	13	7	4173	3064
Airplane (ShapeNet)	11		4045	
Car (FG3D)	20	6	8235	3593
Car (ShapeNet)	18		4472	
Chair (FG3D)	33	20	13054	6781
Chair (ShapeNet)	23		8591	

view aggregation [12], [52], [56], [57], pose estimation [3] have been presented. Parts4Feature [11] integrated a Region Proposal Network (RPN) to detect generally semantic parts in multiple views and then aggregated global shape feature from these generally semantic parts. However, it is still hard for current methods to fully explore the fine-grained details of 3D shapes in the fine-grained classification task. In FG3D-Net, we introduce a hierarchical part-view attention aggregation strategy to extract more discriminative information from generally semantic parts.

III. FG3D DATASET

To evaluate the performance in fine-grained 3D shape classification, we introduce a fine-grained 3D shape (FG3D) dataset. Different from previous datasets such as ShapeNet [1], FG3D aims to evaluate the fine-grained recognition of sub-categories within the same basic category, where 3D shapes may exhibit large intra-subcategory variance and small inter-subcategory variance.

A dataset for fine-grained 3D shape classification needs to fulfill six crucial properties [18], [23]: (1) a large number of shapes for deep networks to capture statistically significant patterns; (2) ground truth labels that enable to quantitatively evaluate the performance in a specific task; (3) unique fine-grained labels for each shape from dozens of subcategories

under the same basic shape category; (4) convenient shape representation as input; (5) 3D file format which deals with the challenges of 3D shape recognition; (6) expandability, i.e., make it easy for the collection to grow over time, to keep the dataset challenging as the performance of learning algorithms improves. To build the FG3D dataset, we collect a large quantity of 3D shape files and classify each 3D shape into one unique subcategory, which strictly follows the above requirements.

Existing datasets are usually composed of 3D shapes from different basic categories such as ShapeNet [9] and ModelNet [35], but they do not satisfy the aforementioned property (3). Although ShapeNet includes multiple subcategories under each basic category, some 3D shapes have multiple or incorrect subcategory labels. This makes it impossible to use ShapeNet for fine-grained 3D shape classification. In addition, the image-based datasets such as FGVC-Aircraft [24] and Car [25] do not satisfy property (5), since they do not contain any 3D shapes.

FG3D is complementary to current datasets by addressing fine-grained classification and satisfying the properties (1)-(6). From the perspective of fine-grained visual classification, FG3D exhibits considerable variation within subcategories, but limited variation among different subcategories under each category. Hence, fine-grained classification on FG3D is a challenging task. To show the differences with ShapeNet,

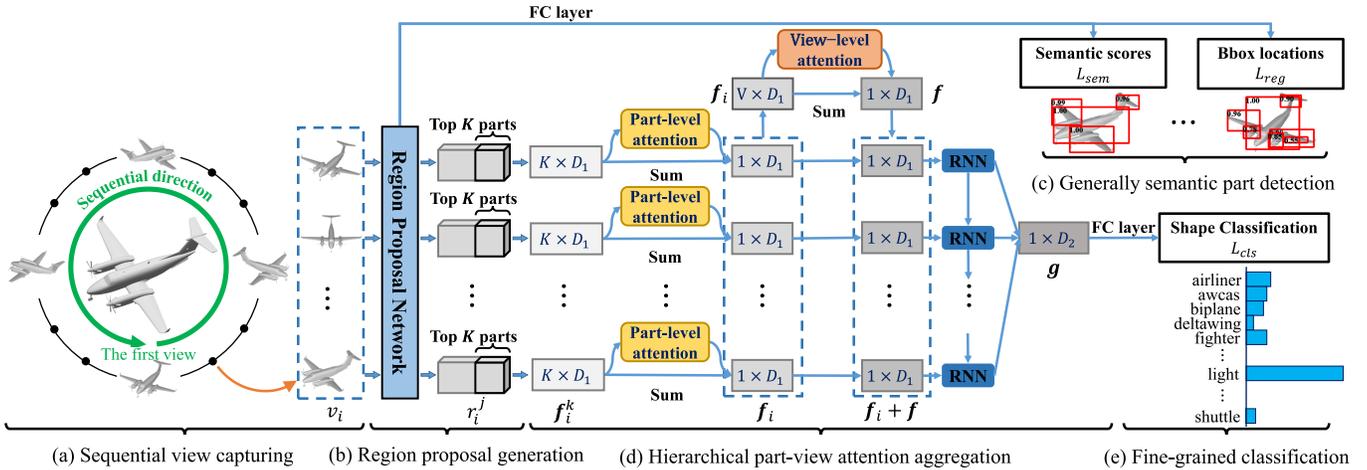


Fig. 3. The framework of FG3D-Net. A sequence of views are first rendered from multiple viewpoints around the input 3D shape in the sequential view capturing module (a). In module (b), all the views are propagated into a region proposal network to generate region proposals and to compute the corresponding proposal features with RoI pooling (b). Then, generally semantic parts in each view are detected by predicting the semantic scores and bounding box (bbox) locations with several FC layers in module (c). Next, in the hierarchical part-view attention aggregation module (d), the top K region proposals according to the semantic scores are selected to extract the global feature for the input 3D shape. There are three different semantic levels in the feature aggregation mechanism, including part-level, view-level, and shape-level. Two bilinear attention mechanisms are integrated to explore the correlation among features in different semantic levels. In addition, an RNN layer is applied to enhance the correlation of view features by taking advantage of the sequential input views. Finally, the global shape representation is extracted by a max-pooling layer, which is applied to the fine-grained classification of 3D shapes in the module (e).

we compare FG3D with ShapeNet for the three categories in FG3D in TABLE II. FG3D contains more 3D objects and subcategories to support the fine-grained classification task.

As shown in TABLE I, FG3D consists of three basic categories including *Airplane*, *Car* and *Chair*, which contain 3,441 shapes in 13 subcategories, 8,235 shapes in 20 subcategories, and 13,054 shapes in 33 subcategories, respectively. We represent each 3D shape by an object format file (.off) with polygonal surface geometry. One can easily convert the .off files into other shape representations, such as rendered views, voxels and point clouds. All shapes in FG3D are collected from multiple online repositories including 3D Warehouse [58], Yobi3D [59] and ShapeNet [9], which contain a massive collection of CAD shapes that are publicly available for research purpose. By collecting 3D shapes over a period of two months, we obtained a collection of more than 20K 3D shapes in three shape categories. We organized these 3D shapes using the WordNet [13] noun “synsets” (synonym sets). WordNet provides a broad and deep taxonomy with over 80K distinct synsets representing distinct noun concepts. This taxonomy has been utilized by ImageNet [5] and ShapeNet [9] to formulate the object subcategories. In our dataset, we also introduce the taxonomy into the collection of 3D shapes, as shown in Fig. 2.

For evaluation, we split the shapes in each categories into training and testing sets. Specifically, the 3D shapes in airplane are split into 3,441 for training and 732 for testing. The cars category contains 7,010 shapes for training and 1,315 shapes for testing. The chairs category contains 11,124 shapes for training and 1,930 shapes for testing.

IV. FG3D-Net

A. Overview

As shown in Fig. 3, The framework of FG3D-Net consists of five main modules including (a) sequential view capturing,

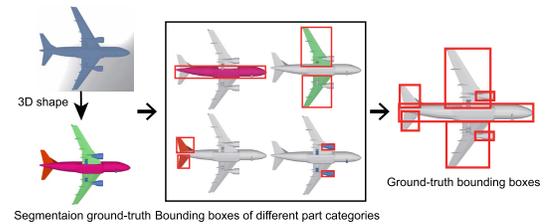


Fig. 4. The generation of ground-truth generally semantic parts from segmentation. According to the segmentation ground-truths, we first individually render each 3D part into 2D views with different colors. Then, we extract the bounding boxes of the colored parts using image processing. Finally, we obtain the ground-truth bounding-boxes of each 3D part in the 2D views.

(b) region proposal generation, (c) generally semantic part detection, (d) hierarchical part-view attention aggregation and (e) fine-grained classification. In particular, the modules (b) and (c) compose a region proposal network (RPN) and cooperate to complete the detection of Generally Semantic Parts (GSPs) from multiple rendered views, which are pre-trained under several part segmentation benchmarks. To construct the region proposal network, we follow the similar strategy as in Parts4Feature [11] which also detects the GSPs from multiple views. By introducing the part information from other segmentation data, our FG3D-Net can integrate the fine-grained details inside local parts.

For each input 3D shape \mathcal{M} from the training set, a view sequence v is first obtained by rendering V views $\{v_i\}$ around \mathcal{M} , such that $v = [v_1, \dots, v_i, \dots, v_V]$ and $i \in [1, V]$, as shown in the module (a) of Fig. 3. Then, a shared convolutional neural network (e.g., VGG19 [60]) abstracts all the views into feature maps of the high-dimensional feature space. By applying a sliding window on the feature maps, numerous region proposals $\{r_i^j\}$ are calculated for each view v_i , where the corresponding proposal features $\{c_i^j\}$ are extracted by a

RoI (Region-of-Interest) pooling layer and $j \in [1, N]$. With proposal features \mathbf{c}_i^j , module (c) learns to predict both the semantic scores and bounding box locations of GSPs within multiple views. Finally, according to the semantic scores, part features $\{\mathbf{c}_i^k\}$ of the top K region proposals $\{r_i^k\}$ in each view v_i of \mathbf{v} are selected for extracting the global shape feature \mathbf{f} of \mathcal{M} with the module (d). The global shape feature \mathbf{f} is propagated through a Fully Connected (FC) layer to provide the classification probability \mathbf{p} for fine-grained 3D shape classification.

B. Generally Semantic Parts (GSPs)

A generally semantic part (GSP) in FG3D-Net indicates a local part in rendered views of a 3D shape. GSPs do not distinguish between different semantic part classes or shape categories, such as the wings of airplanes or the wheels of cars. Different from semantic parts in 3D shape segmentation, a generally semantic part in FG3D-Net represents a local visual part in 2D views, rather than shape parts in 3D. By learning the shape representation from GSPs, our method exploits the fine-grained details of multiple views at three different feature abstraction levels, including part-level, view-level, and shape-level.

In the pre-training stage of GSPs detection, all ground-truth GSPs are generated from several 3D shape segmentation benchmarks. By learning the GSPs from other 3D segmentation datasets, we are able to detect the GSPs in multiple views under our fine-grained classification benchmark without requiring ground-truth GSP supervision in the FG3D dataset. Specifically, we use three 3D shape segmentation benchmarks including ShapeNetCore [61], Labeled-PSB [21], [62], and COSEG [62] to construct the generally semantic part detection benchmark and provide ground-truth GSPs. Here, we follow [63] to split the 3D shapes into training and testing set.

Fig. 4 shows the pipeline to generate the ground-truth bounding boxes for GSPs. Based on the segmentation ground-truth in 3D, we represent each segmented 3D part with a different color. To extract the bounding box for each 3D part, we render each colored 3D part into 2D images. With a simple image processing step, we apply a region property measurement function to calculate the bounding box of the colored parts. To reduce the impact of data noise, we apply a data cleaning step to eliminate the influence of some small parts whose bounding boxes are smaller than 0.45 of the max bounding box in the same part category. In addition, another benefit of using GSPs is eliminating the impact of some incorrect segmentations, where 3D parts may have wrong segmentation labels in the datasets. So far, we have obtained the bounding boxes of the ground truth GSPs within multiple views, where the ground-truth GSPs are applied to train modules (b) and (c) for generally semantic part detection, as shown in Fig. 3.

C. Region Proposal Network

A Region Proposal Network (RPN) takes an image as input and outputs a set of rectangular object proposals, each with an objectness score. As shown in Fig. 3, the modules (a)

and (b) of our FG3D-Net compose a RPN, which is adjusted for detecting GSPs from the multiple views of 3D shapes. Similar to [6], the RPN detects a set of generally semantic part proposals with corresponding semantic scores for each view. Specifically, in module (b), a large number of region proposal candidates $\{r_i^j\}$ are first generated for each view v_i and the corresponding proposal feature $\{\mathbf{c}_i^j\}$ are calculated by a Region-of-Interest (RoI) pooling layer. To complete the region proposal generation, there are a Convolutional Neural Network (CNN) layer, a region proposal calculation layer and a RoI pooling layer in the module (b). Then, in module (c), all the proposal features $\{\mathbf{c}_i^j\}$ are applied to predict both the semantic scores and locations of GSPs with several stacked Fully Connected (FC) layers. According to the predicted semantic scores, the features $\{\mathbf{c}_i^k\}$ of the top K proposals are selected for hierarchically learning the 3D shape representations in module (d).

To review the details of the RPN, a Convolutional Neural Network (CNN) first abstracts multiple input views into the feature maps. The CNN layer is modified from a VGG-19 network proposed in [60], and it produces a feature map \mathbf{c}_i for each view v_i . Secondly, a region proposal layer calculates the region proposals $\{r_i^j\}$ by sliding a small window over the corresponding feature map \mathbf{c}_i . At each sliding-window location, which is centered at each pixel \mathbf{c}_i , a region r_i^j is proposed by regressing its location t_D and predicting a semantic probability p_D with an anchor. The location t_D is a four dimensional vector representing center coordinates, width and height of the part bounding box. All of the above modifications are adjusted from previous Faster-RCNN [64].

To train RPN for predicting the semantic scores p_D of GSPs, we assign a binary label to each region proposal r_i^j to indicate whether r_i^j is a GSP. Specifically, we assign a positive label if the IoU (Intersection-over-Union) overlap between r_i^j and any ground-truth GSP in v_i is higher than a threshold S_D . Note that a single ground-truth box may assign positive labels to multiple anchors; otherwise, we use a negative label. In each view v_i , we apply RoI pooling over region proposal location t_D on feature maps \mathbf{c}_i . Hence, the features $\{\mathbf{c}_i^j\}$ of all N region proposals $\{r_i^j\}$ are high dimensional vectors, which we forward to the generally semantic part detection module as shown in Fig. 3. To reduce the computational cost, we represent each generally semantic part feature in the top K proposals with a high dimensional feature vector \mathbf{f}_i^k that is calculated by $\mathbf{f}_i^k = \text{MAX}(\{\mathbf{c}_i^k\})$. Therefore, we finally select the features $\{\mathbf{f}_i^k\}$ of the top K region proposals $\{r_i^k\}$ according to corresponding semantic scores p_D , which are propagated to the hierarchical part-view attention aggregation module.

With the above definitions, an objective function is employed to optimize the part detection following the multi-task loss in Faster-RCNN [6]. Denote that the ground-truth semantic scores and box locations of positive or negative samples in the RPN are p^* and t^* , respectively. The loss function is formulated as

$$L(p_D, p^*, t_D, t^*) = L_{sem}(p_D, p^*) + \lambda L_{reg}(t_D, t^*), \quad (1)$$

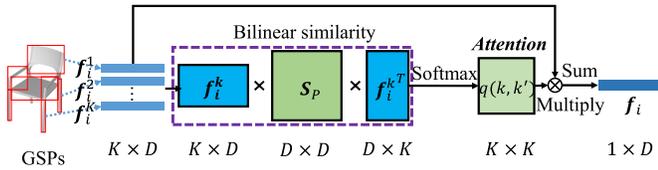


Fig. 5. The demonstration of part-level attention. By calculating the attention values among generally semantic parts, we aggregate K parts features $\{f_i^k\}$ in the same view into a view feature f_i . Firstly, we compute the bilinear similarity score $q(k, k')$ between f_i^k and $f_i^{k'}$ with a softmax layer. Then, we multiply the scores $q(k, k')$ to the input part features to highlight the discriminative features. Here, \otimes indicates multiplication. Finally, a weighted sum is applied to aggregate the highlighted parts features into a view feature f_i .

where L_{sem} measures the accuracy in terms of semantic scores by calculating the cross-entropy loss with part labels, while L_{reg} measures the accuracy of location in terms of the L_1 distance [6]. In the experiment, λ works well with a value of 1. By introducing the architecture of RPN [6], FG3D-Net has the special ability to detect GSPs from multiple views by being pre-trained under other segmentation benchmarks.

D. Hierarchical Part-View Attention Aggregation

In order to extract shape features with fine-grained distinction, the hierarchical part-view attention aggregation module (d) is the key component of FG3D-Net, as shown in Fig. 3, which hierarchically learns the global representation f of the input 3D shape from the features $\{f_i^k\}$ of generally semantic parts. To propagate the features from low-level to high-level, it is important to preserve the detailed information in different levels. Thus, in module (d), there are three different semantic levels including part-level, view-level and shape-level, where the fine-grained relationship among features are explored. To aggregate shape features from low to high, some special designs are introduced including part-level attention, view-level attention and view feature enhancement.

1) *Part-Level Attention*: The target of part-level attention is to aggregate the features of GSPs detected from the same view. As shown in Fig. 5, we first select the top K GSP features $\{f_i^k\}$ of view v_i to extract the corresponding view feature f_i . Traditional approaches such as [1] usually aggregate multiple features by simple max-pooling or mean-pooling operations. However, simple pooling operation suffers from the content information loss within generally semantic parts. As summarized by some previous methods [65]–[68] for fine-grained classification of 2D images, detailed information about local parts usually determines the discriminability of learned object features. To resolve this issue, we propose a bilinear similarity attention to aggregate generally semantic part features $\{f_i^k\}$ into view feature f_i for view v_i in the view sequence v . The part-level attention is designed to take advantage of the relationships among part features to facilitate the feature aggregation procedure. Specifically, a shared matrix S_p is learned to evaluate the mutual correlations among the K generally semantic part features $\{f_i^k\}$ of view v_i . Here the learnable matrix S_p is shared across all views, which aims to capture the general patterns among GSPs in each one

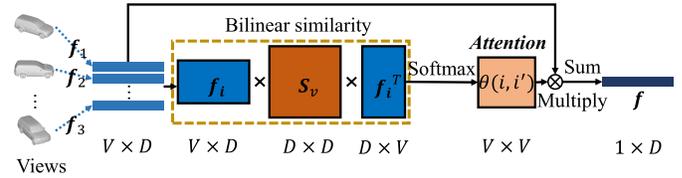


Fig. 6. The demonstration of view-level attention. Similar to the part-level attention, we calculate the bilinear similarity scores $\theta(i, i')$ among the view features from the same 3D object to aggregate view features $\{f_i\}$ into a shape feature f . By applying a multiply operation \otimes and a weighted sum operation, the global feature f is obtained by considering the importance of different view features.

of multiple views. In addition, the learning of S_p can also explore the local patterns of GSPs from different 3D shapes. By applying an attention value to each GSP, the part-level attention can highlight the important GSPs in each view to facilitate feature aggregation.

Given the GSP features $\{f_i^k\}$ of view v_i , the corresponding context of f_i^k is formed by

$$\mathbf{R}_{i,k}^{context} = \{f_i^k\}, \quad k \in [1, K]. \quad (2)$$

For each candidate f_i^k , there is a score $q(k, k')$ measuring the similarity between f_i^k and $f_i^{k'}$ as follows

$$q(k, k') = \frac{\exp(f_i^k S_p f_i^{k'T})}{\sum_{n=1}^K \exp(f_i^k S_p f_i^{nT})}, \quad (3)$$

where S_p is the learnable matrix. All $q(k, k')$ form the bilinear similarity attention matrix with a size of $K \times K$, which represents the correlation among the K GSPs. Without extra selection of highly related ones, the context e_i^k fused for f_i^k is a weighted average over all the candidate parts, as denoted by

$$e_i^k = \sum_{k'=1}^K q(k, k') f_i^{k'}. \quad (4)$$

With the context vector of GSPs in view v_i , the view feature f_i can be calculated by

$$f_i = \sum_{k=1}^K e_i^k, \quad (5)$$

where the view features $\{f_i\}$ are propagated to both the view-level attention and the view feature enhancement module.

2) *View-Level Attention*: In the view-level, we apply a similar strategy to effectively aggregate view features f_i into 3D global features f as depicted in Fig. 6, which contains the global information of an entire 3D shape from multiple views. To learn the attention value of view features from the same 3D object, one learnable parameter matrix S_v is also learned to capture the correlation among views. Therefore, by calculating the bilinear similarity among view features, the attention values are captured, which can also leverage the importance of views in the feature aggregation.

Given the view features $\{f_1, f_2, \dots, f_i, \dots, f_V\}$, a similarity score $\theta(i, i')$ measuring the similarity between f_i and

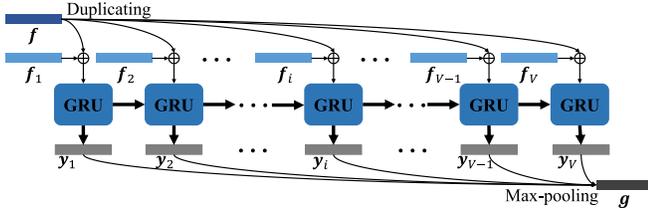


Fig. 7. The demonstration of the view feature enhancement. To enhance the view features, we first add the global shape feature f to f_i by duplicating, where \oplus indicates the sum operation. Then, we adopt a GRU (RNN) to capture the spatial correlation among views. Therefore, we obtain the enhanced view features $\{y_i\}$ which are aggregated into the final global shape representation g by a max-pooling layer.

$f_{i'}$ is computed as

$$\theta(i, i') = \frac{\exp(f_i S_b f_{i'}^T)}{\sum_{l=1}^V \exp(f_i S_b f_l^T)}, \quad (6)$$

where S_b is also a learnable matrix to capture the correlation among views. The learned $\theta(i, i')$ are the bilinear similarity attention matrix with a size of $V \times V$. According to the learned similarity scores, the context \bar{e}_i of view v_i is fused from f_i with a weighted average over the views

$$\bar{e}_i = \sum_{i'=1}^V \theta(i, i') f_{i'}. \quad (7)$$

Finally, the 3D global feature f is computed as

$$f = \sum_{i=1}^V \bar{e}_i. \quad (8)$$

Through the hierarchical part-view attention mechanisms, we obtain the global shape representation f , which can be applied for 3D shape recognition. However, the spatial correlation among sequential views is not fully explored. To take advantage of the prior information of the view sequence v , we additionally utilize a recurrent neural network to enhance current 3D global feature f .

3) *View Feature Enhancement*: An important prior information of sequential input views v is the spatial information among them, where the rendering viewpoints are continuously distributed around the shape in 3D space as shown in Fig. 3. To benefit from the powerful ability of learning from sequential data, FG3D-Net employs a Recurrent Neural Network (RNN) to learn the enhanced view features $\{y_i\}$ from previous view features $\{f_i\}$ and global feature f .

With the view features $\{f_1, f_2, \dots, f_i, \dots, f_V\}$ aggregated from the part level, we first integrate the global shape information by

$$f'_i = f_i + f. \quad (9)$$

The integrated features f'_i form a feature sequence $f'_v = \{f'_1, f'_2, \dots, f'_i, \dots, f'_V\}$. A RNN takes the sequential view features as input and captures the spatial correlation among views. The RNN consists of a hidden state h and an optional output y , which operates on the view feature sequence f'_v . Here each item f'_i is a 512-dimensional feature vector and

the length of f'_v is V which is also the number of steps in the RNN. At each time step $t \in [1, V]$, the hidden state h_t of the RNN is updated by

$$h_t = GRU(h_{t-1}, f'_t), \quad (10)$$

where GRU is a non-linear activation function named gated recurrent unit [69].

A RNN can learn the probability distribution over a sequence by being trained to predict the next item in the sequence. Similarly, at time step t , the output y_t of the RNN can be represented as

$$y_t = W_a h_t, \quad (11)$$

where W_a is a learnable weight matrix. After forwarding the entire input feature sequence, as shown in Fig. 7, the output sequence $\{y_1, y_2, \dots, y_i, \dots, y_V\}$ is acquired, which contains the content information and the spatial information of the entire view feature sequence. To avoid a dependency on the choice of the initial input view feature in the RNN, we learn a more general global shape feature g by adopting a max-pooling layer,

$$g = \max_{i \in [1, V]} \{y_i\}. \quad (12)$$

Through the view feature enhancement, a global representation g is extracted from GSPs in multiple views. Following a FC layer, g is applied to predict 3D shape labels with the cross-entropy loss function, where the softmax function outputs the classification probabilities p . Suppose that there are C shape subcategories of an individual category in the classification, so each probability $p(c)$ can be defined as

$$p(c) = \frac{\exp(W_c g + a_c)}{\sum_{c' \in [1, C]} \exp(W_{c'} g + a_{c'})}, \quad (13)$$

where W and a are the weights of the FC layer to be learned. The objective loss function is the cross-entropy between the predicted probability p and the corresponding ground truth p' ,

$$L(p, p') = - \sum_{c \in [1, C]} p'(c) \log(p(c)). \quad (14)$$

E. Training

There are two different tasks in our FG3D-Net including generally semantic part detection and fine-grained shape classification. To leverage the performance of the two tasks, we use an alternating strategy to train FG3D-Net. First, we train the region proposal network to detect generally semantic parts inside shape views under the processed segmentation benchmarks. Then, while fixing the parameters in the region proposal network, we only update parameters inside the hierarchical part-view attention aggregation branch, as shown in Fig. 1. By repeating the above two steps, we can apply our FG3D-Net to extract discriminative shape features for 3D shapes for fine-grained 3D shape classification. Therefore, the optimization target L_{total} of FG3D-Net consists of a sum of two parts,

$$L_{total} = L(p_D, p^*, t_D, t^*) + \psi L(p, p'), \quad (15)$$

where ψ is a hyperparameter to balance loss terms.

TABLE III
THE KEY PARAMETERS IN FG3D-NET

V	the number of input views
K	the number of GSPs in each view
CT	the cell type of RNN
Dim	the dimension of RNN hidden state
S_D	the threshold of positive GSP
N	the number of proposal candidates
S_p, S_v	the learnable bilinear matrices
λ, ϕ	the ratio of loss function terms

V. EXPERIMENTS

In this section, we conduct comprehensive experiments to validate FG3D-Net for fine-grained classification of 3D shapes. We first explore how some key parameters affect the performance of FG3D-Net and then discuss results of an ablation study to justify our architecture. Finally, under our FD3D dataset, we compare the fine-grained shape classification performance of FG3D-Net with state-of-the-art methods with different 3D shape representations.

A. Network Configuration

As shown in TABLE III, we illustrate the key parameters in our FG3D-Net. In the experiment, we apply default $V = 12$ views of each 3D shape \mathcal{M} , where the viewpoints are uniformly distributed around the shape and each view size is $224 \times 224 \times 3$. After the generally semantic part detection with $S_D = 0.7$, the features of top $K = 20$ parts in each view are propagated to the hierarchical part-view attention aggregation module, as shown in Fig. 3. Following some parameter settings as in Faster RCNN [6], the size of the abstracted feature map c_i is $12 \times 12 \times 512$ for each view. To generate anchors on the feature map, We apply 6 scales and 3 aspect ratios to yield $6 \times 3 = 18$ anchors at each pixel, which ensures a wide coverage of sizes to accommodate region proposals for GSPs that may be partially occluded in some views. The 6 scales relative to the size of the views are [1,2,4,8,16,32], and the 3 aspect ratios are 1 : 1, 1 : 2, and 2 : 1. In general, all anchors lead to $N = 2592 = 12 \times 12 \times 18$ proposal candidates $\{r_i^j\}$ in each view v_i . And each region proposal feature c_i^j extracted by RoI pooling are with size of $512 \times 7 \times 7$. In the part-view attention mechanisms, the learnable bilinear matrices S_p, S_v are with size 512×512 , and the dimensions of all features are 512. In the view feature enhancement module, the GRU cell [69] is adopted in our RNN and the dimension of the hidden state is initialized with 4,096. And the hyperparameters λ, ψ are set to 1 in the loss function. For all experiments, we train our network on a NVIDIA GTX 1,080Ti GPU using ADAM optimizer with an initial learning rate of 0.00001 and a batch size of 1.

B. Parameter Setting

There are several important hyperparameters settings in FG3D-Net. To investigate the influence of these hyperparameters, we performed comparisons with different settings under the category of airplane, which contains 4,173 3D shapes from

TABLE IV
THE EFFECT OF THE VIEW NUMBER V ON THE PERFORMANCE IN FG3D-NET

V	3	6	12	20
Acc (%)	88.39	92.90	93.99	93.99

TABLE V
THE EFFECT OF THE NUMBER OF GENERALLY SEMANTIC PARTS K ON THE PERFORMANCE IN FG3D-NET

K	5	10	20	30	40
Acc (%)	92.90	93.31	93.99	93.58	93.44

TABLE VI
THE EFFECT OF RNN CELL TYPES (CTs) ON THE PERFORMANCE IN FG3D-NET

CT	BasicRNN	LSTM	BidirectionRNN	GRU
Acc (%)	91.80	92.76	93.44	93.99

13 subcategories. We first explore the number of input views V , which determines the coverage of input 3D shapes from different view angles. In this experiment, we keep all other hyperparameters fixed and modify the number of rendered views V from 3 to 20. All views are obtained by rendering around each 3D shape as shown in MVCNN [1]. The results are listed in TABLE IV, which shows the accuracy trend with increasing the number of views. With $V = 12$ or $V = 20$ views as input, FG3D-Net reaches the similar instance accuracy of 93.99%. This is because 12 views have already covered most of the details of 3D shapes. In addition, 12 views have been widely used as input of view-based methods, such as MVCNN [1], SeqViews2SeqLabels [12], 3D2SeqViews [57] and Parts4Feature [11]. For fair comparison with these methods and leveraging the computational complexity of the whole network, we also adopt 12 views as input in our FG3D-Net.

In the following experiments, we keep the number of input views $V = 12$. In TABLE V, we study the effect of the number K of generally semantic parts in the hierarchical part-view attention aggregation module in Fig. 3 (d), where K ranges from 5 to 40. The best accuracy 93.99% is reached at $K = 20$, which can achieve a better coverage of the fine-grained details of 3D shapes.

To investigate the effect of the RNN cell type (CT) in our *view feature enhancement* step, we show the results with different RNN cells in TABLE VI. We observe that the GRU cell outperforms other RNN cells such as BasicRNN, LSTM and BidirectionRNN. In particular, in the BidirectionRNN cell, there are two GRU cells of different directions.

Moreover, we explore the effect of the RNN's hidden state dimension (Dim) which affects the learning ability of the recurrent neural network. As depicted in TABLE VII, the dimension of GRU cells is modified from 512 to 5,120. The best performance is achieved at $Dim = 4096$, which can better capture the spatial correlation of sequential views.

TABLE VII

THE EFFECT OF THE DIMENSION (DIM) OF GRU CELLS ON THE PERFORMANCE

Dim	512	1024	2048	4096	5120
Acc (%)	93.03	93.44	93.17	93.99	92.76

TABLE VIII

THE EFFECT OF THE ATTENTION MECHANISM ON THE PERFORMANCE UNDER AIRPLANE IN FG3D-NET

Metric	OVA	OPA	NA	NR
Acc (%)	93.31	93.58	92.90	91.94

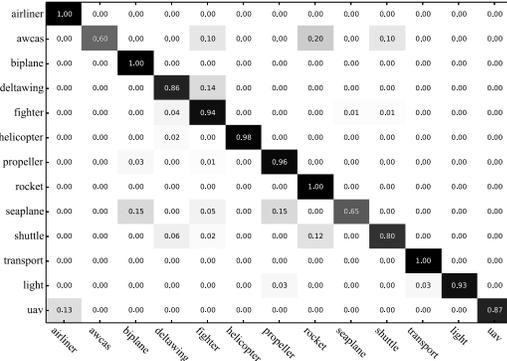


Fig. 8. The classification confusion matrix under Airplane.

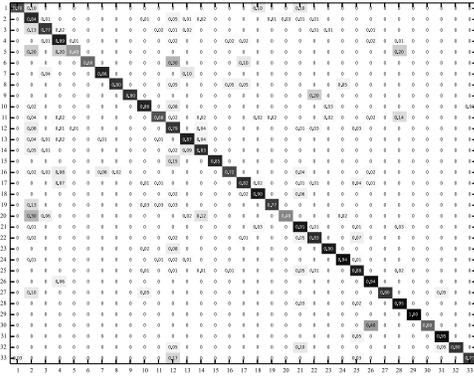


Fig. 9. The classification confusion matrix under Chair.

C. Ablation Study

In order to reveal the effect of the novel elements in FG3D-Net, such as attention mechanisms at different semantic levels, we performed ablation studies to justify their effectiveness. We evaluate the performance of FG3D-Net with only part-level attention (OPA), only view-level attention (OVA), no attention (NA) or no RNN (NR). Specifically, when we remove the attention mechanism, we set all attention values to the same constant. For example, we set the part-level attention value to $\frac{1}{20}$ in OVA and the view-level attention value $\frac{1}{12}$ in OPA. TABLE VIII illustrates the effectiveness of our attention mechanisms in learning highly discriminative representations for 3D shapes. The results show that both part attention and view attention play an important role in extracting fine-grained details from multiple views. Without the RNN layer,

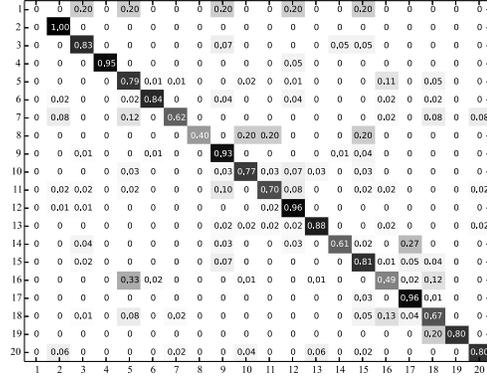


Fig. 10. The classification confusion matrix under Car.

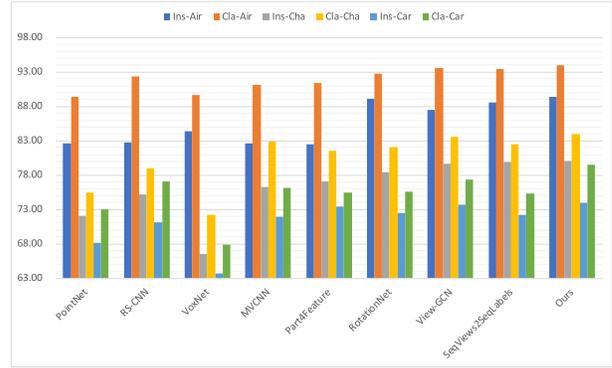


Fig. 11. The intuitive display of classification accuracies under the FG3D Dataset, including instance accuracies (Ins-Air, Ins-Cha, Ins-Car) and class accuracies (Cla-Air, Cla-Cha, Cla-Car).

the performance of FG3D-Net drops significantly. This is mainly caused by the reduction of network parameters and the lack of spatial correlations among views. In other words, the spatial correlation among views is important for the fine-grained classification of 3D shapes.

As for the components of neural networks, there is a Region Proposal Network in the GSP detection procedure, which contains a VGG-19 network, a RoI pooling layer and four FC layers. And for the subsequent networks, we select 20 GSPs from each view of total 12 views. The GSPs detected from multiple views play an important role in capturing the fine-grained details of 3D shapes, but it also increases the complexity of our method. To better demonstrate the effectiveness of our method, we have evaluated the computational complexity of FG3D-Net on a NVIDIA GTX 1,080Ti GPU. Specifically, the model size and the average forward time are 832Mb and 358.19ms, respectively, where our FG3D-Net adopts 12 views of a 3D shape as input.

D. Fine-Grained Visual Classification

We carry out the experiments in the fine-grained classification of 3D shapes under the proposed FG3D dataset with three categories including *Airplane*, *Car* and *Chair*. As shown in TABLE I, there are 4,173, 8,325 and 13,054 shapes in the categories, which are split into training and testing sets. To evaluate the performance of FG3D-Net, we compare our method with several state-of-the-art 3D shape

TABLE IX

WE COMPARE FINE-GRAINED SHAPE CLASSIFICATION ACCURACY(%) FOR AIRPLANE, CHAIR AND CAR CATEGORIES IN OUR FG3D DATASET

Method	Modality	Input	Airplane		Chair		Car	
			Class	Instance	Class	Instance	Class	Instance
PointNet [42]	point	1024×3	82.67	89.34	72.07	75.44	68.12	73.00
PointNet++ [43]	point	1024×3	87.26	92.21	78.11	80.78	70.30	75.21
SO-Net [70]	point	1024×3	66.10	82.92	62.78	70.05	53.38	59.32
Point2Sequence [71]	point	1024×3	87.52	92.76	74.90	79.12	64.67	73.54
DGCNN [72]	point	1024×3	88.38	93.17	71.74	77.51	65.27	73.61
RS-CNN [73]	point	1024×3	82.81	92.35	75.14	78.96	71.17	77.11
VoxNet [74]	voxel	32^3	84.45	89.62	66.50	72.18	63.75	67.83
MVCNN [1]	view	12×224^2	82.57	91.11	76.27	82.90	71.88	76.12
RotationNet [3]	view	12×224^2	89.11	92.76	78.45	82.07	72.53	75.59
3D2SeqViews [57]	view	12×224^2	89.41	92.21	76.26	77.10	63.85	66.39
3DViewGraph [75]	view	12×224^2	88.21	93.85	79.88	83.58	71.65	77.11
View-GCN [76]	view	12×224^2	87.44	93.58	79.70	83.63	73.68	77.34
SeqViews2SeqLabels [12]	view	12×224^2	88.52	93.44	79.89	82.54	72.23	75.36
Parts4Feature [11]	view	12×224^2	82.55	91.39	77.08	81.61	73.42	75.44
Ours	view	12×224^2	89.44	93.99	80.04	83.94	74.03	79.47

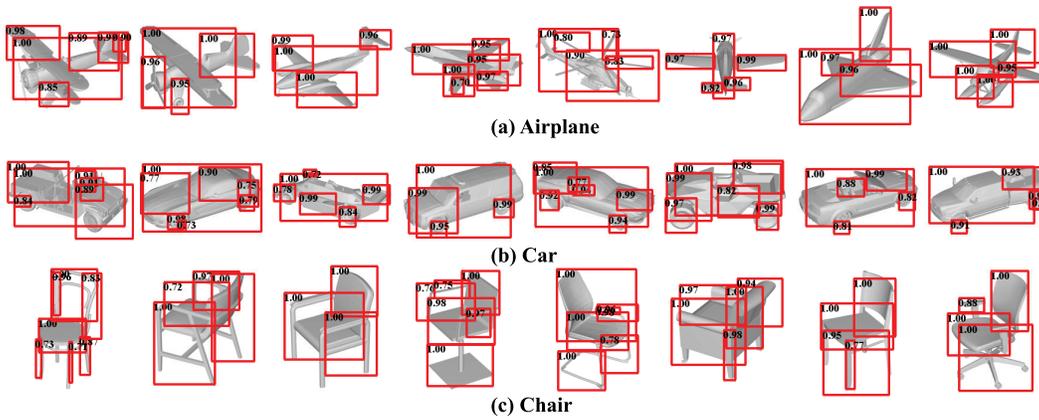
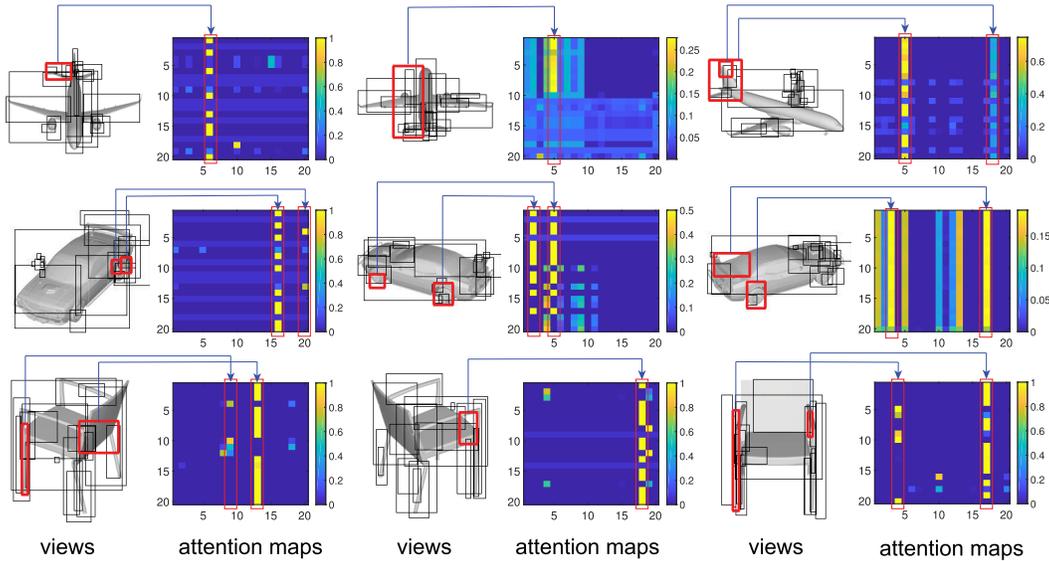


Fig. 12. The results of generally semantic part detection, where the semantic score of each part is larger than 0.8.

Fig. 13. Visualization of the part-level attention mechanism. There are 9 views from the 3 categories of our FG3D dataset, where the top 20 GSPs are drawn in each image. The corresponding bilinear similarity attention map $q(k, k')$ among GSPs in each view with a size of 20×20 is shown on the right of the view. We highlight the discriminative GSPs with red bounding boxes and use arrows to indicate the corresponding columns on the attention maps.

classification methods, which are trained under different 3D shape representations including point clouds, rendered views, and 3D voxels. In TABLE IX, we conduct the numeri-

cal comparisons including PointNet [42], PointNet++ [43], Point2Sequence [71], DGCNN [72], RS-CNN [73], MVCNN [1], SeqViews2SeqLabels [52], RotationNet [3], View-GCN

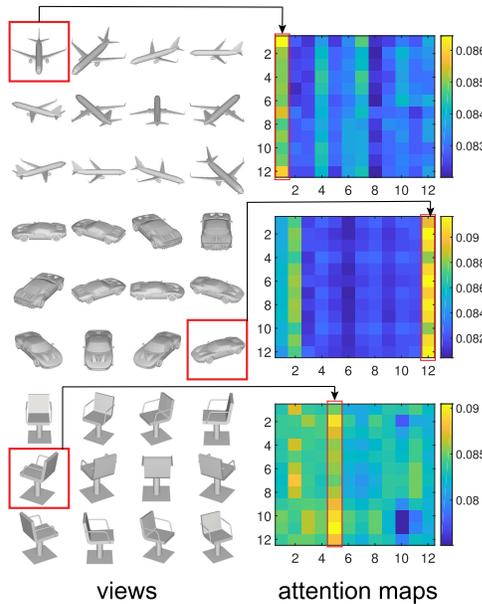


Fig. 14. Visualization of the view-level attention mechanism. There are three 3D shapes selected from three different categories in our FG3D dataset. In each row, we show the 12 views of each 3D shape on the left and the attention map $\theta(i, i')$ of the views with a size of 12×12 on the right. We highlight the discriminative views with red bounding boxes and use arrows to indicate the corresponding columns on the attention maps.

[76] and Parts4Feature [11]. For these methods, we reproduce network structures using the public source code and evaluate them under our FG3D dataset.

In the comparison with other methods, we adopt MVCNN [1] as our baseline. MVCNN is the pioneering study in 3D shape recognition from multiple views, which has achieved satisfactory performance in shape classification. MVCNN leverages 12 views rendered around each 3D shape as input. We follow this camera setting and evaluate the performance of other view-based methods, including RotationNet [3], 3D2SeqViews [57], 3DViewGraph [75], View-GCN [76], SeqViews2SeqLabels [12] and Parts4Feature [11]. RotationNet learns the best camera setting to capture the view-specific feature representation for 3D shapes, which has obtained superior performance to previous state-of-the-art methods under ModelNet40 and ModelNet10. Recent View-GCN [76] also achieves high performances via view-based graph convolution network under ModelNet benchmarks. Therefore, we have compared FG3D-Net with RotationNet and View-GCN under the FG3D dataset using the same view sequence obtained from each 3D shape in this work. To utilize the sequential views, SeqViews2SeqLabels adopts a recurrent neural network to aggregate multiple views of 3D shapes. However, SeqViews2SeqLabels cannot explore the fine-grained information among GSPs extracted from the views. Parts4Feature also detects generally semantic parts inside views similar to FG3D-Net. However, Parts4Feature cannot fully explore the fine-grained correlation among features at different semantic levels. In particular, there is no view-level feature in Part4Feature, which results in a lack of view-level information in feature aggregation.

In addition, we include methods that take other data formats as input. For point-based methods, we choose the pioneering PointNet [42] and PointNet++ [43] as the comparison targets. We further include state-of-the-art methods such as SO-Net [70] and Point2Sequence [71]. To translate our FG3D dataset into point clouds, we apply Poisson Disk Sampling [77] to obtain 1,024 points for each 3D shape. All point-based methods are trained with 1,024 points as input under FG3D. VoxNet [74], also listed in TABLE IX, is the pioneering work in voxel-based methods for 3D shape recognition.

As shown in TABLE IX, our FG3D-Net outperforms other state-of-the-art methods and achieves the highest classification accuracies in all three categories. The results suggest that FG3D-Net can take advantage of generally semantic part detection to integrate fine-grained details in multiple views. With a hierarchical aggregation strategy, we fully explore the correlation of features at different semantic levels.

To better demonstrate our classification results, we visualize the confusion matrix of our classification result under Airplane, Chair and Car in Fig. 9, Fig. 8 and Fig. 10, respectively. In each confusion matrix, an element in the diagonal line means the classification accuracy in a class, while other elements in the same row means the misclassification accuracy. The large diagonal elements show that FG3D-Net is good at classifying large-scale 3D shapes. And to show the classification accuracies more intuitively, we draw the accuracy bar of FG3D-Net under different basic categories as illustrated in Fig. 11. From the accuracy bar, FG3D-Net achieves higher performances than other compared methods.

VI. VISUALIZATION

In this section, we visualize some important properties of our FG3D-Net. Firstly, we show examples of detected GSPs under our FG3D testing set. In Fig. 12, we draw the bounding box of GSPs whose semantic score is larger than 0.8. We observe that FG3D-Net can extract the bounding boxes of discriminative GSPs, which supports representation learning of 3D shapes.

There are two attention mechanisms in FG3D-Net, including part-level and view-level attention. To learn the global representation of 3D shapes, these attention mechanisms are important to preserve the fine-grained details inside GSPs from multiple views. To intuitively show the effectiveness of these attention mechanisms, we draw some samples of part-level attention in Fig. 13 and view-level attention in Fig. 14, respectively. In Fig. 13, we show the attention map $q(k, k')$ of 9 views from 3 different categories. In each view, we draw the bounding boxes of top 20 GSPs, where the most discriminative GSPs with large attention values are in red. We use arrows to indicate the correspondence of GSPs and attention values in the attention map $q(k, k')$, where each column represents the attention value of each GSP. Similar to the part-level attention, we visualize the view-level attention in Fig. 14, which shows the 12 views and the corresponding view attention map $\theta(i, i')$ of a 3D shape in each row. We also use arrows to indicate the correspondence of views and attention values on the attention maps. The visualization results show that both part-level

attention and view-level attention are effective to capture the fine-grained information in the feature aggregation.

VII. CONCLUSION AND FUTURE WORK

We proposed FG3D-Net, a novel model to learn 3D global features via hierarchical feature aggregation from GSPs. To evaluate the performance of FG3D-Net, we introduced a new fine-grained 3D shape classification dataset. In the existing methods, the fine-grained details of generally semantic parts and the correlation of features in different semantic levels are usually ignored, which limits the discriminability of learned 3D global features. To resolve these disadvantages, FG3D-Net employs a region proposal neural network to detect GSPs from multiple views, which identifies discriminative local parts inside views to learn the fine-grained details of 3D shapes. In addition, we leverage part-level and view-level attention mechanisms to effectively aggregate features in different semantic levels, which utilize the correlations among features to highlight the discriminative features. Finally, a recurrent neural network is adopted to capture the spatial information among views from multiple viewpoints, which takes advantage of prior information about the sequential input views. Experimental results show that our method outperforms the state-of-the-art under the proposed 3D fine-grained dataset.

Although FG3D-Net learns 3D shape global features from GSPs to achieve high performance in fine-grained shape classification, it still suffers from three limitations. First, FG3D-Net can only detect the limited types of GSPs, since the number of 3D shapes in the existing segmentation benchmarks is still limited. For example, the performance of our FG3D-Net may be further improved by integrating the fine-grained parts from PartNet [23]. Second, the number of categories in our FG3D dataset is somewhat small, where only three categories are currently included. Thus, FG3D-Net merely performs well under the current input setting and on a small number of shape categories, even with the help of bilinear similarity attention and RNN. Third, there is no validation set in the splitting of our FG3D dataset. A good validation set can help to obtain better hyperparameters in the network. In the future, it is worth to explore unsupervised methods to detect GSPs inside multiple views and to further extend our FG3D dataset, including adding more categories and making a partition of validation set.

REFERENCES

- [1] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller, "Multi-view convolutional neural networks for 3D shape recognition," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 945–953.
- [2] S. Bai, X. Bai, Z. Zhou, Z. Zhang, Q. Tian, and L. J. Latecki, "GIFT: Towards scalable 3D shape retrieval," *IEEE Trans. Multimedia*, vol. 19, no. 6, pp. 1257–1271, Jun. 2017.
- [3] A. Kanazaki, Y. Matsushita, and Y. Nishida, "RotationNet: Joint object categorization and pose estimation using multiviews from unsupervised viewpoints," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5010–5019.
- [4] Z. Han, M. Shang, X. Wang, Y.-S. Liu, and M. Zwicker, "Y2Seq2Seq: Cross-modal representation learning for 3D shape and text by joint reconstruction and prediction of view and word sequences," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 126–133.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [6] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [7] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The Caltech-UCSD birds-200-2011 dataset," *Comput. Neural Syst., California Inst. Technol., Pasadena, CA, USA, Tech. Rep. CNS-TR-2011-001*, 2011.
- [8] A. Khosla, N. Jayadevaprakash, B. Yao, and F.-F. Li, "Novel dataset for fine-grained image categorization: Stanford dogs," in *Proc. CVPR Workshop Fine-Grained Vis. Categorization*, vol. 2, no. 1, 2011.
- [9] A. X. Chang *et al.*, "ShapeNet: An information-rich 3D model repository," 2015, *arXiv:1512.03012*. [Online]. Available: <http://arxiv.org/abs/1512.03012>
- [10] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser, "The princeton shape benchmark," in *Proc. IEEE Shape Modeling Appl.*, 2004, pp. 167–178.
- [11] Z. Han, X. Liu, Y.-S. Liu, and M. Zwicker, "Parts4Feature: Learning 3D global features from generally semantic parts in multiple views," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Aug. 2019, pp. 766–773.
- [12] Z. Han *et al.*, "SeqViews2SeqLabels: Learning 3D global features via aggregating sequential views by RNN with attention," *IEEE Trans. Image Process.*, vol. 28, no. 2, pp. 658–672, Feb. 2019.
- [13] G. A. Miller, "WordNet: A lexical database for English," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, 1995.
- [14] S. Seok, C. D. Onal, K.-J. Cho, R. J. Wood, D. Rus, and S. Kim, "Meshworm: A peristaltic soft robot with antagonistic nickel titanium coil actuators," *IEEE/ASME Trans. Mechatronics*, vol. 18, no. 5, pp. 1485–1497, Oct. 2013.
- [15] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang, "Pixel2Mesh: Generating 3D mesh models from single RGB images," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 52–67.
- [16] F. Bogo, J. Romero, M. Loper, and M. J. Black, "FAUST: Dataset and evaluation for 3D mesh registration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 3794–3801.
- [17] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, *Numerical Geometry of Non-Rigid Shapes*. Cham, Switzerland: Springer, 2008.
- [18] S. Koch *et al.*, "ABC: A big CAD model dataset for geometric deep learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 9601–9611.
- [19] H. Benhabiles, J.-P. Vandeboer, G. Lavoue, and M. Daoudi, "A framework for the objective evaluation of segmentation algorithms using a ground-truth of human segmented 3D-models," in *Proc. IEEE Int. Conf. Shape Model. Appl.*, Jun. 2009, pp. 36–43.
- [20] R. Hu, L. Fan, and L. Liu, "Co-segmentation of 3D shapes via subspace clustering," *Comput. Graph. Forum*, vol. 31, no. 5, pp. 1703–1713, 2012.
- [21] E. Kalogerakis, A. Hertzmann, and K. Singh, "Learning 3D mesh segmentation and labeling," *ACM Trans. Graph.*, vol. 29, no. 4, p. 102, Jul. 2010.
- [22] Y. Wang, S. Asafi, O. Van Kaick, H. Zhang, D. Cohen-Or, and B. Chen, "Active co-analysis of a set of shapes," *ACM Trans. Graph.*, vol. 31, no. 6, p. 165, 2012.
- [23] K. Mo *et al.*, "PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 909–918.
- [24] S. Maji, E. Rahtu, J. Kannala, M. Blaschko, and A. Vedaldi, "Fine-grained visual classification of aircraft," 2013, *arXiv:1306.5151*. [Online]. Available: <http://arxiv.org/abs/1306.5151>
- [25] J. Krause, M. Stark, J. Deng, and L. Fei-Fei, "3D object representations for fine-grained categorization," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, Sydney, NSW, Australia, Dec. 2013, pp. 554–561.
- [26] D. Wang, Z. Shen, J. Shao, W. Zhang, X. Xue, and Z. Zhang, "Multiple granularity descriptors for fine-grained categorization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 2399–2406.
- [27] T.-Y. Lin, A. RoyChowdhury, and S. Maji, "Bilinear CNN models for fine-grained visual recognition," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1449–1457.
- [28] M. Jaderberg *et al.*, "Spatial transformer networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 2017–2025.
- [29] P. Sermanet, A. Frome, and E. Real, "Attention for fine-grained categorization," in *Proc. Int. Conf. Learn. Represent. Workshop*, 2015.
- [30] N. Zhang, J. Donahue, R. Girshick, and T. Darrell, "Part-based R-CNNs for fine-grained category detection," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2014, pp. 834–849.
- [31] M. Simon and E. Rodner, "Neural activation constellations: Unsupervised part model discovery with convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1143–1151.

- [32] Y. Zhang *et al.*, “Weakly supervised fine-grained categorization with part-based image representation,” *IEEE Trans. Image Process.*, vol. 25, no. 4, pp. 1713–1725, Apr. 2016.
- [33] Z. Han, Z. Liu, J. Han, C.-M. Vong, S. Bu, and X. Li, “Unsupervised 3D local feature learning by circle convolutional restricted Boltzmann machine,” *IEEE Trans. Image Process.*, vol. 25, no. 11, pp. 5331–5344, Nov. 2016.
- [34] Z. Han, Z. Liu, J. Han, C.-M. Vong, S. Bu, and C. L. P. Chen, “Mesh convolutional restricted Boltzmann machines for unsupervised learning of features with structure preservation on 3-D meshes,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 10, pp. 2268–2281, Oct. 2017.
- [35] Z. Wu *et al.*, “3D ShapeNets: A deep representation for volumetric shapes,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1912–1920.
- [36] Z. Han, G. Qiao, Y.-S. Liu, and M. Zwicker, “SeqXY2SeqZ: Structure learning for 3D shapes by sequentially predicting 1D occupancy segments from 2D coordinates,” in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 607–625.
- [37] Z. Han, B. Ma, Y.-S. Liu, and M. Zwicker, “Reconstructing 3D shapes from multiple sketches using direct shape optimization,” *IEEE Trans. Image Process.*, vol. 29, pp. 8721–8734, Sep. 2020.
- [38] A. Sharma, O. Grau, and M. Fritz, “VConv-DAE: Deep volumetric shape learning without object labels,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Cham, Switzerland: Springer, 2016, pp. 236–250.
- [39] C. R. Qi, H. Su, M. NieBner, A. Dai, M. Yan, and L. J. Guibas, “Volumetric and multi-view CNNs for object classification on 3D data,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 5648–5656.
- [40] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum, “Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 82–90.
- [41] S. Muralikrishnan, V. G. Kim, and S. Chaudhuri, “Tags2Parts: Discovering semantic regions from shape tags,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2926–2935.
- [42] R. Q. Charles, H. Su, M. Kaichun, and L. J. Guibas, “PointNet: Deep learning on point sets for 3D classification and segmentation,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 652–660.
- [43] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “PointNet++: Deep hierarchical feature learning on point sets in a metric space,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5099–5108.
- [44] X. Wen, Z. Han, X. Liu, and Y.-S. Liu, “Point2SpatialCapsule: Aggregating features and spatial relationships of local regions on point clouds using spatial-aware capsules,” *IEEE Trans. Image Process.*, vol. 29, pp. 8855–8869, 2020.
- [45] X. Wen, Z. Han, G. Youk, and Y.-S. Liu, “CF-SIS: Semantic-instance segmentation of 3D point clouds by context fusion with self-attention,” in *Proc. 28th ACM Int. Conf. Multimedia*, Oct. 2020, pp. 1661–1669.
- [46] X. Wen, T. Li, Z. Han, and Y.-S. Liu, “Point cloud completion by skip-attention network with hierarchical folding,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 1939–1948.
- [47] Z. Han, C. Chen, Y.-S. Liu, and M. Zwicker, “DRWR: A differentiable renderer without rendering for unsupervised 3D structure learning from silhouette images,” in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 3994–4005.
- [48] Z. Han, C. Chen, Y.-S. Liu, and M. Zwicker, “ShapeCaptioner: Generative caption network for 3D shapes by learning a mapping from parts detected in multiple views to sentences,” in *Proc. 28th ACM Int. Conf. Multimedia*, Oct. 2020, pp. 1018–1027.
- [49] Y. Yang, C. Feng, Y. Shen, and D. Tian, “FoldingNet: Point cloud auto-encoder via deep grid deformation,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 206–215.
- [50] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas, “Learning representations and generative models for 3D point clouds,” in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 40–49.
- [51] Z. Han, X. Wang, Y.-S. Liu, and M. Zwicker, “Multi-angle point cloud-VAE: Unsupervised feature learning for 3D point clouds from multiple angles by joint self-reconstruction and half-to-half prediction,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 10441–10450.
- [52] Z. Han, M. Shang, Y.-S. Liu, and M. Zwicker, “View inter-prediction GAN: Unsupervised representation learning for 3D shapes by learning global shape memories to support local view predictions,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 8376–8384.
- [53] B. Shi, S. Bai, Z. Zhou, and X. Bai, “DeepPano: Deep panoramic representation for 3-D shape recognition,” *IEEE Signal Process. Lett.*, vol. 22, no. 12, pp. 2339–2343, Dec. 2015.
- [54] K. Sfikas, T. Theoharis, and I. Pratikakis, “Exploiting the PANORAMA representation for convolutional neural network classification and retrieval,” in *Proc. 3DOR*, vol. 6, 2017, p. 7.
- [55] E. Johns, S. Leutenegger, and A. J. Davison, “Pairwise decomposition of image sequences for active multi-view recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3813–3822.
- [56] C. Wang, M. Pelillo, and K. Siddiqi, “Dominant set clustering and pooling for multi-view 3D object recognition,” in *Proc. Brit. Mach. Vis. Conf.*, 2017, pp. 1–12.
- [57] Z. Han *et al.*, “3D2SeqViews: Aggregating sequential views for 3D global feature learning by CNN with hierarchical attention aggregation,” *IEEE Trans. Image Process.*, vol. 28, no. 8, pp. 3986–3999, Aug. 2019.
- [58] C. Goldfeder and P. Allen, “Autotagging to improve text search for 3D models,” in *Proc. IEEE Int. Conf. Shape Model. Appl.*, Jun. 2008, pp. 355–358.
- [59] *Yobi3D-Free 3D Model Search Engine*. Accessed: Dec. 11, 2019. [Online]. Available: <https://www.yobi3d.com>
- [60] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *Proc. Int. Conf. Learn. Represent.*, 2015, pp. 1–14.
- [61] L. Yi *et al.*, “A scalable active framework for region annotation in 3D shape collections,” *ACM Trans. Graph.*, vol. 35, no. 6, p. 210, Nov. 2016.
- [62] X. Chen, A. Golovinskiy, and T. Funkhouser, “A benchmark for 3D mesh segmentation,” *ACM Trans. Graph.*, vol. 28, no. 3, p. 73, 2009.
- [63] E. Kalogerakis, M. Averkiou, S. Maji, and S. Chaudhuri, “3D shape segmentation with projective convolutional networks,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3779–3788.
- [64] R. Girshick, J. Donahue, T. Darrell, and J. Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 580–587.
- [65] H. Zheng, J. Fu, T. Mei, and J. Luo, “Learning multi-attention convolutional neural network for fine-grained image recognition,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5209–5217.
- [66] Z. Wang, T. Chen, G. Li, R. Xu, and L. Lin, “Multi-label image recognition by recurrently discovering attentional regions,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 464–472.
- [67] P. Guo and R. Farrell, “Aligned to the object, not to the image: A unified pose-aligned representation for fine-grained recognition,” in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2019, pp. 1876–1885.
- [68] W. Ge, X. Lin, and Y. Yu, “Weakly supervised complementary parts models for fine-grained image classification from the bottom up,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 3034–3043.
- [69] K. Cho *et al.*, “Learning phrase representations using RNN encoder-decoder for statistical machine translation,” in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2014, pp. 1724–1734.
- [70] J. Li, B. M. Chen, and G. H. Lee, “SO-net: Self-organizing network for point cloud analysis,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 9397–9406.
- [71] X. Liu, Z. Han, Y.-S. Liu, and M. Zwicker, “Point2sequence: Learning the shape representation of 3D point clouds with an attention-based sequence to sequence network,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 8778–8785.
- [72] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, “Dynamic graph CNN for learning on point clouds,” *ACM Trans. Graph.*, vol. 38, no. 5, pp. 1–12, Nov. 2019.
- [73] Y. Liu, B. Fan, S. Xiang, and C. Pan, “Relation-shape convolutional neural network for point cloud analysis,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 8895–8904.
- [74] D. Maturana and S. Scherer, “VoxNet: A 3D convolutional neural network for real-time object recognition,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2015, pp. 922–928.
- [75] Z. Han, X. Wang, C. M. Vong, Y.-S. Liu, M. Zwicker, and C. L. P. Chen, “3DViewGraph: Learning global features for 3D shapes from a graph of unordered views with attention,” in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Aug. 2019, pp. 758–765.
- [76] X. Wei, R. Yu, and J. Sun, “View-GCN: View-based graph convolutional network for 3D shape analysis,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 1850–1859.
- [77] C. Feng, *Cython Wrapper of Poisson Disk Sampling of a Triangle Mesh in VCGLIB*. Accessed: Jan. 12, 2020. [Online]. Available: <https://github.com/simbafirst/trimesh2pointcloud>



Xinhai Liu received the B.S. degree in computer science and technology from the Huazhong University of Science and Technology, China, in 2017. He is currently pursuing the Ph.D. degree with the School of Software, Tsinghua University. His research interests include deep learning, 3D shape analysis, and 3D pattern recognition.



Yu-Shen Liu (Member, IEEE) received the B.S. degree in mathematics from Jilin University, China, in 2000, and the Ph.D. degree from the Department of Computer Science and Technology, Tsinghua University, Beijing, China, in 2006. From 2006 to 2009, he was a Postdoctoral Researcher with Purdue University. He is currently an Associate Professor with the School of Software, Tsinghua University. His research interests include shape analysis, pattern recognition, machine learning, and semantic search.



Zhizhong Han received the Ph.D. degree from Northwestern Polytechnical University, China, in 2017. He is currently a Postdoctoral Researcher with the Department of Computer Science, University of Maryland, College Park, MD, USA. He is also a Research Member with the BIM Group, Tsinghua University, China. His research interests include machine learning, pattern recognition, feature learning, and digital geometry processing.



Matthias Zwicker (Associate Member, IEEE) received the Ph.D. degree from ETH in Zurich, Switzerland, in 2003. He is currently a Professor with the Department of Computer Science, University of Maryland, College Park, MD, USA, where he holds the Reginald Allan Hahne Endowed E-nnovate Chair. Before joining the University of Maryland, he was an Assistant Professor with the University of California, San Diego, CA, USA, and a Professor with the University of Bern, Switzerland. His research interests include computer graphics for high-quality rendering, point-based methods for rendering and modeling, 3D geometry processing, and data-driven modeling and animation.