

ShapeCaptioner: Generative Caption Network for 3D Shapes by Learning a Mapping from Parts Detected in Multiple Views to Sentences

Zhizhong Han

Tsinghua University and University of Maryland
College Park, Maryland
h312h@umd.edu

Yu-Shen Liu*

School of Software, BNRist, Tsinghua University
Beijing, China
liuyushen@tsinghua.edu.cn

Chao Chen

School of Software, BNRist, Tsinghua University
Beijing, China
thss15_chenc@163.com

Matthias Zwicker

University of Maryland
College Park, Maryland
zwicker@cs.umd.edu

ABSTRACT

3D shape captioning is a challenging application in 3D shape understanding. Captions from recent multi-view based methods reveal that they cannot capture part-level characteristics of 3D shapes. This leads to a lack of detailed part-level description in captions, which human tend to focus on. To resolve this issue, we propose *ShapeCaptioner*, a generative caption network, to perform 3D shape captioning from semantic parts detected in multiple views. Our novelty lies in learning the knowledge of part detection in multiple views from 3D shape segmentations and transferring this knowledge to facilitate learning the mapping from 3D shapes to sentences. Specifically, ShapeCaptioner aggregates the parts detected in multiple colored views using our novel part class specific aggregation to represent a 3D shape, and then, employs a sequence to sequence model to generate the caption. Our outperforming results show that ShapeCaptioner can learn 3D shape features with more detailed part characteristics to facilitate better 3D shape captioning than previous work.

CCS CONCEPTS

• Computing methodologies → Computer vision tasks.

KEYWORDS

3D shape captioning, Description generation, Semantic part detection, Multiple views, RNN.

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '20, October 12–16, 2020, Seattle, WA, USA.

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-7988-5/20/10...\$15.00
<https://doi.org/10.1145/3394171.3413889>



Figure 1: Humans tend to describe 3D shapes by focusing on their semantic parts (in red) with their various attributes (in blue).

ACM Reference Format:

Zhizhong Han, Chao Chen, Yu-Shen Liu, and Matthias Zwicker. 2020. ShapeCaptioner: Generative Caption Network for 3D Shapes by Learning a Mapping from Parts Detected in Multiple Views to Sentences. In *28th ACM International Conference on Multimedia (MM '20)*, October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3394171.3413889>

1 INTRODUCTION

Jointly understanding 3D shapes and sentences is an important challenge in 3D shape analysis. For example, generated captions are helpful for visually impaired people to understand what a 3D shape looks like, including the category, color, form, and material of the 3D shape, as shown in the examples in Fig. 1. This motivates us to address the issue of automatically generating captions for 3D shapes.

Recently, Text2Shape [1] made an important contribution by proposing a 3D-Text dataset, where voxel-based 3D shapes and their corresponding captions are paired together. To reduce voxel complexity by representing 3D shapes as view sequences, Y²Seq2Seq [16] was presented to jointly learn a bilateral mapping between view sequences and word sequences. Although Y²Seq2Seq can produce plausible captions for 3D shapes, its ability to generate detailed description for parts is limited because the model does not capture part-level characteristics. However, human-provided captions often focus on part-level details, as the manually annotated ground truth examples in Fig. 1 illustrate. Therefore, how to absorb part characteristics in 3D shape captioning remains challenging.

To resolve this issue, we propose a novel deep neural network, *ShapeCaptioner*, to generate captions for 3D shapes. By representing a 3D shape as a view sequence, ShapeCaptioner

aims to learn a mapping from parts detected in the view sequence to a caption describing the 3D shape. Specifically, ShapeCaptioner first leverages 3D shape segmentation benchmarks to learn the knowledge of detecting parts in terms of their geometry from multiple views, and then, transfers this knowledge to incrementally learn the ability of detecting parts with various attributes from multiple colored views in 3D-Text dataset. ShapeCaptioner further employs a novel part class specific aggregation to aggregate the detected parts over all views for the captioning of a 3D shape. The part class specific aggregation can represent a 3D shape by capturing more part characteristics from different views, which facilitates more detailed 3D shape captioning. In summary, our contributions are as follows:

- i) We propose ShapeCaptioner to enable 3D shape captioning from semantic parts detected in multi-views, which facilitates more detailed 3D shape captioning.
- ii) We introduce a method to learn the knowledge of part detection in multi-views from 3D shape segmentation benchmarks to facilitate 3D shape captioning.
- iii) With our novel part class specific aggregation, we effectively capture part characteristics to represent 3D shapes for better shape captioning, inspired by the way humans describe 3D shapes in terms of semantic parts.

2 RELATED WORK

Image captioning. There is a large volume of work on image captioning. Here we review methods based on object detection, which are most similar to our work. Currently, end-to-end deep learning approaches [26, 27] are most effective for image captioning. These methods try to learn to generate captions from global image features. However, using a global feature limits their interpretability. Other work [4, 24, 43, 46, 48, 49] employs object-level semantics to generate captions. These methods represent images based on occurring semantic concepts or objects. In [48], explicitly detected objects are employed with their category, size and layout to generate captions. In contrast, based on a bag of word model, impressive captions can also be generated by only using the explicitly detected objects [43].

Our method is different from these methods in four aspects. First, the parts we want to detect are more various than the objects. Second, part detection becomes more challenging when considering multiple unaligned and varying viewpoints. Third, how to aggregate the detected parts over different views to enable learning the mapping from 3D shapes to captions represents an additional problem. Fourth, a final obstacle is that there is no labeled dataset available to learn the knowledge of part detection in multiple views of 3D shapes. Yet ShapeCaptioner can resolve these issues to generate captions with detailed part characteristics.

3D shape captioning. Although deep learning based 3D shape understanding has made significant progress [6–21, 23, 29–31, 33, 34, 44, 45, 47?], 3D shape captioning has been less explored due to the lack of training dataset. However, the recently proposed 3D-Text dataset [1] has enabled the

research in this area. Y²Seq2Seq [16] employs multiple views as a 3D shape representation to address the cubic complexity of voxel representations. It learns 3D shape features by aggregating the global feature of each view. Although Y²Seq2Seq can generate plausible captions for 3D shapes, it often fails to generate captions with local part details. To resolve this issue, ShapeCaptioner represents a 3D shape as a set of parts detected in multiple views, which not only avoids the cubic complexity of voxels but also enables the ability of capturing part characteristics. This leads to captions that are more similar to the manually annotated ground truth, which usually includes part details, such as color, form, material, and texture of parts shown in Fig. 1.

3 OVERVIEW

ShapeCaptioner learns to generate a caption \mathbf{t} for a 3D shape \mathbf{s} under a 3D-Text dataset \mathbf{D} which provides ground truth (\mathbf{s}, \mathbf{t}) pairs, where shape \mathbf{s} is represented by a 128^3 dimensional voxel cube and contains various attributes, such as color, material and texture, as demonstrated in Fig. 1. ShapeCaptioner represents shape \mathbf{s} as a colored view sequence $\mathbf{v} = \{v_i, i \in [1, V]\}$, and detects important parts from each view v_i to capture part characteristics for caption generation. However, dataset \mathbf{D} does not have ground truth parts in multiple views for ShapeCaptioner to learn from.

To resolve this problem, we first leverage a separate 3D shape segmentation benchmark \mathbf{B} to learn the knowledge of part geometry detection in multiple views, i.e., detecting parts only in terms of geometry rather than attributes, as demonstrated in Fig. 2 (a) and (b). This is because there is no color, material or texture available in benchmark \mathbf{B} . Then, we transfer this knowledge to 3D-Text dataset \mathbf{D} to establish the ground truth parts in multiple colored views which enables ShapeCaptioner to incrementally learn the ability of detecting parts in terms of both geometry and attributes, as demonstrated in Fig. 2 (b), (c) and (d). Finally, ShapeCaptioner learns a mapping from parts detected in multiple views of shape \mathbf{s} to its description \mathbf{t} to facilitate 3D shape captioning, as demonstrated in Fig. 2 (d), (e) and (f).

4 SHAPECAPTIONER

Part geometry detection. As shown in Fig. 2 (a), we propose a method to obtain part geometry ground truth in multiple views from 3D shape segmentation benchmark \mathbf{B} . We use the 3D shape segmentation benchmarks involved in [25] as segmentation ground truth.

As illustrated in Fig. 3, starting from each 3D mesh in benchmark \mathbf{B} , our method first voxelizes the 3D mesh into voxels, along with labelling each voxel according to the segmentation ground truth on the mesh. In the voxelization, we randomly sample 100 points on each triangle face of the mesh, and label each sampled point by the label of the triangle face. This enables us to perform label voting among points located in the same voxel in the labelling of each voxel. Using randomly sampled points could resist the imbalance effect of triangle face size. Then, we render the voxelized shape

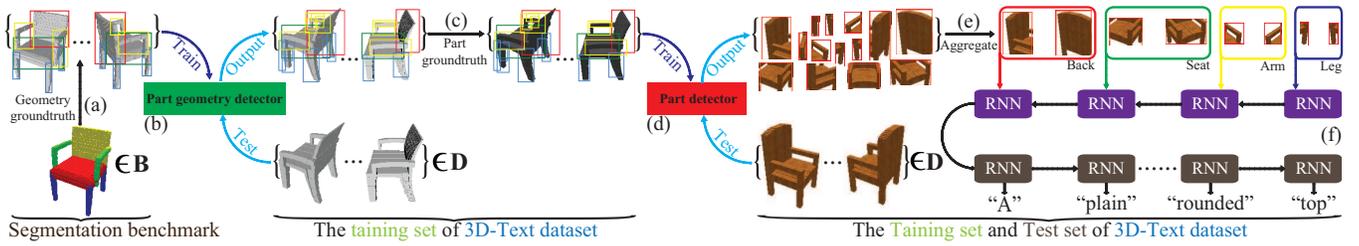


Figure 2: The demonstration of ShapeCaptioner. ShapeCaptioner first learns the knowledge of detecting part geometry in multiple views from shape segmentation benchmark B ((a) and (b)). By transferring this knowledge to 3D-Text dataset D, ShapeCaptioner learns the ability of detecting parts containing both geometry and attributes ((b), (c) and (d)). Finally, ShapeCaptioner generates captions by learning the mapping from the detected parts to the description ((d), (e) and (f)).

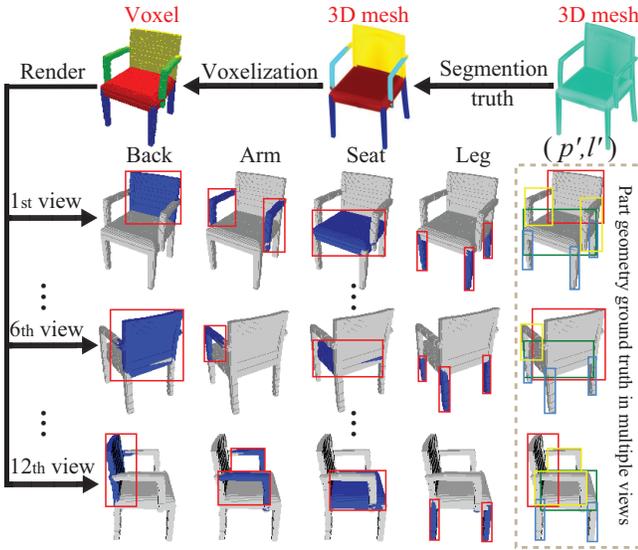


Figure 3: Generating part geometry ground truth.

from 12 viewpoints and locate different part geometries in each view. Specifically, from each viewpoint, we separately highlight voxels belonging to the same part class in blue, and compute the bounding box of each blue region. We denote p' as a one hot probability distribution to indicate which part class the bounding box belongs to, and denote l' as the location of the bounding box, where (p', l') forms the part geometry ground truth.

Take the first viewpoint in Fig. 3 for example, we separately obtain the bounding box of the region formed by voxels in back, arm, seat and leg class of a chair, and finally, we obtain the part geometry ground truth in the first view. By repeating this process, we obtain the part geometry ground truth (p', l') in multiple views, as shown by the bounding box on shapes in the dashed box in Fig. 3.

The reason why we obtain views with part geometry ground truth by rendering the voxelized shapes in benchmark B is to avoid the domain gap. Because 3D shapes in 3D-Text dataset D are voxelized, there would be a domain gap if we learn

from part geometry ground truth in views rendered from meshes, while detecting part geometry from views rendered from these voxelized shapes.

Subsequently, we train a fasterRCNN [35] as a part geometry detector G under the obtained part geometry ground truth (p', l') , as shown in Fig. 2 (b). This enables ShapeCaptioner to detect part geometries (p_g, l_g) from any view v by minimizing the objective function below,

$$O_G(p_g, p', l_g, l') = O_p(p_g, p') + \lambda O_l(l_g, l'), \quad (1)$$

where O_p measures the accuracy in terms of probability by the cross-entropy function of part class labels, while O_l measures the accuracy in terms of location by the robust L_1 function as in [5]. The parameter λ balances O_p and O_l , and a value of 1 works well in all our experiments.

Parts detection. Part geometry detection is not enough for ShapeCaptioner to generate captions with detailed part characteristics, since the detected geometries lack various attributes, such as color or texture. Hence, it also needs to be able to detect parts with attributes.

To resolve this issue, we transfer the learned knowledge of part geometry detection to the 3D-Text dataset D. We introduce a bounding box mapping method to establish part ground truth which contains both geometry and attributes for ShapeCaptioner to learn from.

As demonstrated in Fig. 2 (b), in the test stage of part geometry detector G , we first render multiple views from each 3D shape in the training set of dataset D without color, and then, employ the trained part geometry detector G to detect part geometries (p_g, l_g) in these views. This process is further illustrated by the first column in Fig. 4. Subsequently, we render the 3D shape again but with color, and map the detected part geometries (p_g, l_g) to the corresponding views with color. Finally, we regard the mapped part geometries in views with color as the part ground truth (p'', l'') , where $l'' = l_g$ and p'' is a one hot probability distribution by setting the entry of $\text{argmax}(p_g)$ to be 1. Here, we map the detected part geometries with $p_g > 0.7$.

Finally, we further push the part geometry detector G to detect parts (p, l) from multiple views with color by fine-tuning G under the part ground truth (p'', l'') . We rename G

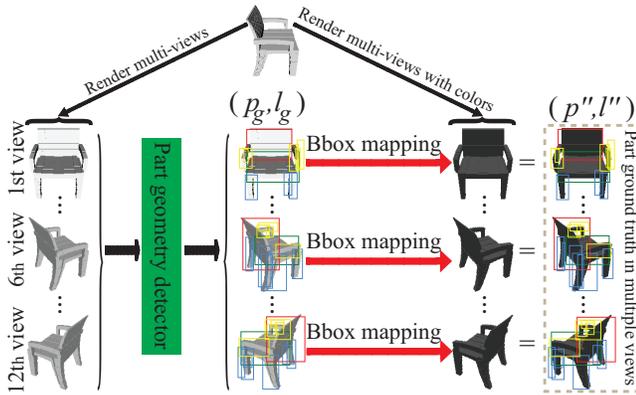


Figure 4: Generating part ground truth.

as part detector R , as shown in Fig. 2 (d). Note that these part ground truth (p'', l'') are only obtained from shapes in the training set of 3D-Text dataset \mathbf{D} . After training, R can detect parts from multiple colored views of shapes in the test set of dataset \mathbf{D} . Similar to Eq. (1), we minimize the following objective function with a λ of 1,

$$O_R(p, p'', l, l'') = O_p(p, p'') + \lambda O_l(l, l''). \quad (2)$$

Part class specific aggregation. Given a 3D shape \mathbf{s} , ShapeCaptioner generates its caption \mathbf{t} from the parts detected in multiple colored views v_i of \mathbf{s} , as demonstrated in Fig. 2 (e). We denote the j -th part detected by R from v_i as (p_i^j, l_i^j) . We select the detected parts with $p_i^j > \rho$ to represent the 3D shape \mathbf{s} , where ρ is a probability threshold.

ShapeCaptioner aggregates the selected parts $\{(p_i^j, l_i^j)\}$ over all V views in terms of different part classes to represent shape \mathbf{s} , as demonstrated in Fig. 2 (e). For each one of C part classes, we group the parts (p_i^j, l_i^j) belonging to the same c -th part class if $\text{argmax}(p_i^j)$ is c . Then, we obtain a part class specific feature \mathbf{F}_c by aggregating all parts in the c -th part class using pooling procedure as follows,

$$\mathbf{F}_c = \text{pool}_{\text{argmax}(p_i^j)=c}(\mathbf{f}_i^j), i \in [1, V], \quad (3)$$

where \mathbf{f}_i^j is a 4096 dimensional feature extracted from the fc7 layer of part detector R . Finally, we use \mathbf{F} to represent shape \mathbf{s} as a sequence of part class specific feature \mathbf{F}_c as follows, where $c \in [1, C]$,

$$\mathbf{F} = [\mathbf{F}_1, \dots, \mathbf{F}_c, \dots, \mathbf{F}_C]. \quad (4)$$

The reason we represent a 3D shape \mathbf{s} as a combination of part class specific features \mathbf{F}_c is to preserve as much part characteristics as possible in the aggregation process. This could reduce the impact of one part class on the others in the multi-view scenario, which enables \mathbf{F}_c to comprehensively describe what the semantic part looks like over all V views of shape \mathbf{s} .

Captioning from parts. ShapeCaptioner leverages the sequence \mathbf{F} of part class specific features \mathbf{F}_c from \mathbf{s} to generate a sequence of words t_n as a caption \mathbf{t} , where $\mathbf{t} =$

$[t_1, \dots, t_n, \dots, t_N]$ and $n \in [1, N]$. We cast this problem into a sequence to sequence translation model, and implement this seq2seq model by a RNN encoder and a RNN decoder, as demonstrated in Fig. 2 (f). The RNN encoder encodes \mathbf{F} by inputting \mathbf{F}_c at each one of C steps, while the RNN decoder dynamically decodes each word t_n in \mathbf{t} . Thus, ShapeCaptioner learns the mapping from 3D shape to sentences by minimizing the following objective function,

$$O_{\mathbf{t}} = - \sum_{t_n \in \mathbf{t}} \log p(t_n | t_{<n}, \mathbf{F}), \quad (5)$$

where t_n is the n -th word in the word sequence \mathbf{t} , $t_{<n}$ represents the words in front of t_n , $p(t_n | t_{<n}, \mathbf{F})$ is the probability of correctly predicting the n -th word according to the previous words $t_{<n}$ and the 3D shape feature \mathbf{F} . Note that the optimization is conducted under the training set of 3D-Text dataset \mathbf{D} . After training, we can generate captions for shapes in the test set of dataset \mathbf{D} .

5 EXPERIMENTS AND ANALYSIS

We present the exploration on the parameters, ablation studies and the comparison with state-of-the-art methods.

Dataset and metrics. We evaluate ShapeCaptioner under the 3D-Text cross-modal dataset [1], which consists of a primitive subset and a ShapeNet subset. We only employ the ShapeNet subset, because the 3D shapes in the primitive subset are too simple to extract parts. The ShapeNet subset contains 15,038 shapes and 75,344 descriptions in only two classes, i.e., chairs and tables. We employ the same training/test splitting in each shape class as [1, 16]. Specifically, the chair class is formed by 5954 training shapes and 641 test shapes while the table class contains 7592 training shapes and 851 test shapes. In addition, we employ 487 chairs (out of 537) and 481 tables (out of 520) involved in segmentation benchmarks [25] to train part geometry detector G , respectively, which avoids the 3D shapes that are also in the test set of 3D-Text dataset.

We employ BLUE [32], CIDEr [40], METEOR [3], and ROUGE [28] to evaluate the quality of generated captions according to the ground truth captions, where these metrics are abbreviated as “B-1”, “B-2”, “B-3”, “B-4”, and “C”, “M”, “R”, respectively in the following tables.

Initialization. We extract 3587 unique words from captions of 3D shapes in the training set of ShapeNet subset to form the vocabulary for caption generation. Each word is represented by a 512 dimensional embedding, which is learned along with the other involved parameters in training. Both RNN encoder and decoder are implemented by GRU cells [2]. The learning rate in all the experiments is 0.00001.

In addition, the number of part classes is $C = 4$ in the chair class, while $C = 3$ in the table class, and ShapeCaptioner is trained under each class respectively. We employ $\rho = 0.8$ to select parts detected in $V = 12$ views to represent a 3D shape, and then, use max pooling to aggregate all the selected parts over V views.



Figure 5: Comparison between captions generated by different methods, where $\rho > 0.8$.

Table 1: The comparison on H . $\rho = 0.8$, $V = 12$.

H	B-1	B-2	B-3	B-4	M	R	C
16	0.122	0.072	0.041	0.021	0.121	0.168	0.002
32	0.937	0.917	0.894	0.878	0.550	0.847	1.789
64	0.794	0.639	0.537	0.475	0.284	0.551	0.644
128	0.761	0.604	0.493	0.421	0.273	0.529	0.560
256	0.772	0.615	0.503	0.431	0.276	0.534	0.576
512	0.736	0.559	0.442	0.371	0.265	0.503	0.506

Parameters. Here we compare some important parameters under the chair class.

We first explore the effect of dimension H of the RNN encoder and decoder by comparing $H \in \{16, 32, 64, 128, 256, 512\}$. As shown in Table 1, the performance is increased with increasing H until $H = 32$, and then, degenerates gradually when H becomes larger. We believe this is caused by over-fitting because the training data is not large enough. In addition, we compare the generated captions under different H in Fig. 6 (a), where the detected parts for caption generation are also briefly shown on the views. We also observe the gradually degenerated captions when H becomes larger. For example, the color description becomes inaccurate in the caption with $H = 128$. Also, some descriptions are repeating in the captions with $H = 256$ and $H = 512$. In the following, we set H to 32.

Then, we explore the probability threshold ρ of selecting parts to represent a 3D shape. We compare $\rho \in \{0.6, 0.7, 0.8, 0.9\}$. There would be more selected parts when ρ is smaller while the quality of the selected parts is lower, and vice versa, as demonstrated in Fig. 6 (b). We believe both the number and the quality would affect the discriminability of 3D shape features. This is because smaller number of parts would decrease the ability of resisting the effect of viewpoint changing while lower quality would contain inaccurate part characteristics.

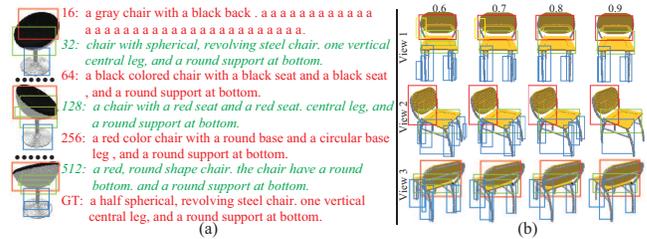


Figure 6: The comparison on (a) H and (b) ρ .

Table 2: The comparison on ρ . $H = 32$, $V = 12$.

ρ	B-1	B-2	B-3	B-4	M	R	C
0.6	0.225	0.140	0.085	0.053	0.154	0.261	0.046
0.7	0.715	0.576	0.476	0.415	0.260	0.515	0.523
0.8	0.937	0.917	0.894	0.878	0.550	0.847	1.789
0.9	0.463	0.409	0.383	0.366	0.235	0.366	0.428

In Table 2, we found $\rho = 0.8$ performs best, hence we use this setting in the following.

Finally, we want to know how the number V of views affects the performance. In this experiment, we use the first $\{1, 4, 8, 12\}$ views to select parts, and then, learn to generate captions. As shown in Table 3, the performance keeps improving along with the increasing number of views. This comparison shows that more views would provide more part characteristics to learn from, which also decreases the effect of inaccurately detected parts. This can also be observed in the comparison of the generated captions under different V in Fig. 7, where the parts employed for caption generation are also briefly shown. For example, we cannot get plausible captions when too few views are available, such as $V = 1$. The material and the color gradually appear in the captions when

Table 3: The comparison on V . $H = 32, \rho = 0.8$.

V	B-1	B-2	B-3	B-4	M	R	C
1	0.324	0.210	0.137	0.097	0.158	0.320	0.061
4	0.777	0.638	0.540	0.478	0.292	0.567	0.724
8	0.876	0.779	0.707	0.656	0.374	0.674	1.108
12	0.937	0.917	0.894	0.878	0.550	0.847	1.789

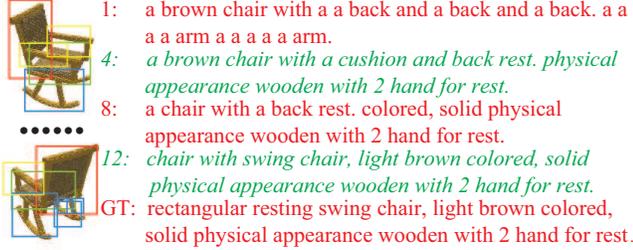


Figure 7: The comparison on V .

Table 4: Part aggregation. $H = 32, \rho = 0.8, V = 12$.

	B-1	B-2	B-3	B-4	M	R	C
Mean	0.220	0.133	0.077	0.044	0.148	0.253	0.018
Mixed	0.843	0.744	0.668	0.617	0.360	0.653	1.019
Max(L)	0.722	0.541	0.433	0.374	0.239	0.482	0.422
MaxAll	0.714	0.593	0.520	0.476	0.299	0.535	0.694
Max	0.937	0.917	0.894	0.878	0.550	0.847	1.789

V increases from 4 to 12. We do not explore the results with more views because of the limited computational capacity. In the following, we use $V = 12$.

Ablation studies. We first highlight our part class specific aggregation for representing a 3D shape. In the former experiments, we employ max pooling to obtain the part class specific aggregation F_c . Here, we try mean pooling to compare with the results of max pooling. As shown by the result of “Mean” in Table. 4, we find that mean pooling is not as good as max pooling (“Max”) to aggregate parts in multiple views. To further justify this point, we conduct another experiment to combine max pooling and mean pooling together. Specifically, we use max pooling to aggregate parts in the same part class in the same view while further using mean pooling to obtain F_c by aggregating the same part class over different views. As shown by the result of “Mixed”, although it is better than the result of “Mean”, it is still worse than the result of “Max”. In addition, we also highlight the idea of part class specific features. As shown by the result of “Max-All”, we max pool all parts over views into a single feature while ignoring the part class, and leverage this feature to generate captions. The degenerated results show that parts in different classes would affect each other in the aggregation. In addition, we compare our employed GRU with LSTM cell in RNN. As shown by “Max(L)”, GRU cell is more suitable in our problem. The captions in Fig. 8 also show the similar comparison results.



Figure 8: Comparison on part aggregation.

Table 5: Representations for chairs. $H = 32, \rho = 0.8, V = 12$.

	B-1	B-2	B-3	B-4	M	R	C
RNN	0.231	0.148	0.090	0.055	0.172	0.252	0.009
CNN	0.337	0.218	0.173	0.152	0.153	0.285	0.221
VMax	0.409	0.244	0.181	0.157	0.176	0.307	0.229
VMean	0.494	0.338	0.251	0.214	0.209	0.381	0.301
Part	0.937	0.917	0.894	0.878	0.550	0.847	1.789

Subsequently, we highlight the advantage of our part based features over view based features. Similar to [16], we employ a VGG19 [37] to extract the feature of each view, and employ a RNN encoder to aggregate these view features for caption generation. As shown in Table. 5, without the multi-task scenario and constraints in [16], this aggregation by RNN (“RNN”) cannot obtain satisfactory results. In addition, we also try a CNN+RNN architecture similar to the method of image captioning [42]. We employ the first view of each shape to caption a 3D shape. Although the results of “CNN” are a little bit better than “RNN”, there is still room to improve. We further try to pool the $V = 12$ view features together into a single feature, and convey this feature to the RNN decoder for caption generation. We find both mean pooling (“VMean”) and max pooling (“VMax”) works well on aggregating views to generate captions, and mean pooling is better than max pooling. However, all these view based shape features cannot capture part characteristics to generate better captions than our part based features (“Part”). Moreover, we also observe similar results in the table class in Table 6.

In addition, we elaborate on the results in Table 5 and Table 6 by cumulative distribution in Fig. 11. For each metric, we use 11 values as probes, and calculate the percentage of samples over the whole test set whose metric scores are bigger than each probe, respectively. These comparisons also demonstrate our significant improvement by higher percentage.

Visualization of detected parts. ShapeCaptioner employs an effective way of detecting parts from multiple colored views of a 3D shape. As demonstrated by the consistent part detection results in Fig. 9, where the bounding boxes of the detected parts are shown on all the $V = 12$ views, ShapeCaptioner can detect reasonable parts for caption generation by understanding the complex geometry of semantic parts without the impacts by viewpoints and colors.

Table 6: Representations for tables. $H = 32, \rho = 0.8, V = 12$.

	B-1	B-2	B-3	B-4	M	R	C
RNN	0.310	0.194	0.110	0.054	0.161	0.317	0.073
CNN	0.451	0.279	0.186	0.143	0.190	0.343	0.228
VMax	0.450	0.285	0.210	0.180	0.192	0.344	0.283
VMean	0.530	0.367	0.274	0.229	0.225	0.414	0.385
Part	0.860	0.755	0.675	0.620	0.362	0.664	1.099

Table 7: Comparison with others. $H = 32, \rho = 0.8, V = 12$.

Method	B-1	B-2	B-3	B-4	M	R	C
ShowAndTell [42]	0.494	0.338	0.251	0.214	0.209	0.381	0.301
Video2Text [41]	0.670	0.430	0.260	0.150	0.210	0.450	0.270
GIFT2Text [39]	0.610	0.350	0.210	0.120	0.160	0.360	0.140
SLR [36]	0.400	0.170	0.080	0.040	0.110	0.240	0.050
Y2S [16]	0.800	0.650	0.540	0.460	0.300	0.560	0.720
Ours(C)	0.937	0.917	0.894	0.878	0.550	0.847	1.789
Ours(T)	0.860	0.755	0.675	0.620	0.362	0.664	1.099
Ours	0.899	0.836	0.785	0.749	0.456	0.756	1.444

Comparison with other methods. We evaluate ShapeCaptioner by comparing it with the state-of-the-art methods. As shown in Table 7, we compare to the methods which are able to generate captions from images or image sequences, such as ShowAndTell [42] for image captioning, Video2Text [41] for video captioning, GIFT2Text [39] for GIF understanding, SLR [36] for video understanding, and Y²Seq2Seq (“Y2S”) [16] for 3D shape understanding. The results of SandT are produced with mean pooling of all views as shape feature for caption generation, while the results of GIF2T and SLR are produced by the nearest retrieval in the joint feature space of shape and caption.

We can see that ShapeCaptioner (“Ours”) significantly outperforms the other view-based methods in all metrics, where “Our” is the average of results under chair (“Ours(C)”) and table (“Ours(T)”) classes. We believe our results benefit from the ability of understanding parts of 3D shapes, which captures more part characteristics to generate better captions in a more similar way to humans. We also obtain the captions in the test set generated by Y²Seq2Seq, and compare Y²Seq2Seq and SandT (“VMean”) to ShapeCaptioner in Fig. 11 in all metrics, where we also observe significant improvement.

Real image test. We further evaluate ShapeCaptioner under a real image set in Fig. 10. What we want to show here is that ShapeCaptioner can also help to caption images by leveraging the detected parts, although it is only trained under 3D data. We select real chair and table images from the Stanford Online Products dataset [38], and we use the trained ShapeCaptioner (“Ours(C)”) and “Ours(T)”) in Table 7 to conduct the results. We first compare our method to the state-of-the-art image captioning methods including CaptionBot [22] and NeuralTalk2 [26] in Fig. 10. We find ShapeCaptioner can generate more detailed and accurate descriptions for parts in captions. This benefits from the detected semantic parts with part details, which also justifies

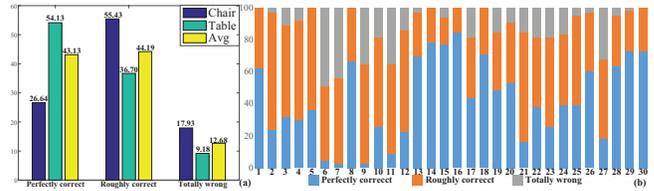


Figure 12: (a) The statistical analysis of the user study over all cases. (b) The statistical analysis of each case.

that ShapeCaptioner has the ability to overcome the gap between rendered views and real images. Note that the generated captions lack variety due to the absence of training on the real images and captions.

To further evaluate the generated captions, we conduct a user study (URL) over randomly selected 30 chairs and tables. We provide all 66 participants each one of the 30 cases accompanied with the generated captions. Then, we ask the participants to select one of three options for each case, i.e., perfectly correct (All descriptions about parts are correct), roughly correct (Some descriptions about parts are correct), and totally wrong (All descriptions about parts are wrong), to evaluate how well the generated caption matches the shape in the real image. Finally, we show the statistical results in Fig. 12 (a). We can see that participants gave almost 90% cases (“Avg”) perfectly and roughly correct, where the results on each case are elaborated in Fig. 12 (b). In addition, our results in the table class (“Table”) are better than the results in the chair class (“Chair”) in terms of percentage of perfectly correct. This is because tables are usually simpler than chairs, and there are more training samples in the table class than the chair class in the 3D-Text dataset.

6 CONCLUSION

We propose ShapeCaptioner to better caption a 3D shape by leveraging more part characteristics. ShapeCaptioner successfully learns the ability of semantic part detection from multiple views of 3D shapes in segmentation benchmarks, and effectively transfers this ability to 3D-Text dataset to caption 3D shapes from the parts detected in multiple colored views. Moreover, our part class specific aggregation can also preserve the part characteristics from different views. Our outperforming results indicate that captioning from parts can produce more accurate descriptions for parts, which is also more similar to human’s way of describing 3D shapes.

ACKNOWLEDGMENTS

This work was supported by National Key R&D Program of China (2018YFB0505400), in part by the NSFC (61672307), National Key R&D Program of China (2019YFB1405703) and TC190A4DA/3, Tsinghua-Kuaishou Institute of Future Media Data, and NSF (award 1813583).



Figure 9: Demonstration of part detection under 3D-Text dataset, where $\rho > 0.8$.



Figure 10: Comparison with CaptionBot (Black) [22] and NeuralTalk2 (Blue) [26], our results (red) presents more detailed and accurate part characteristics under real images, where $\rho > 0.5$.

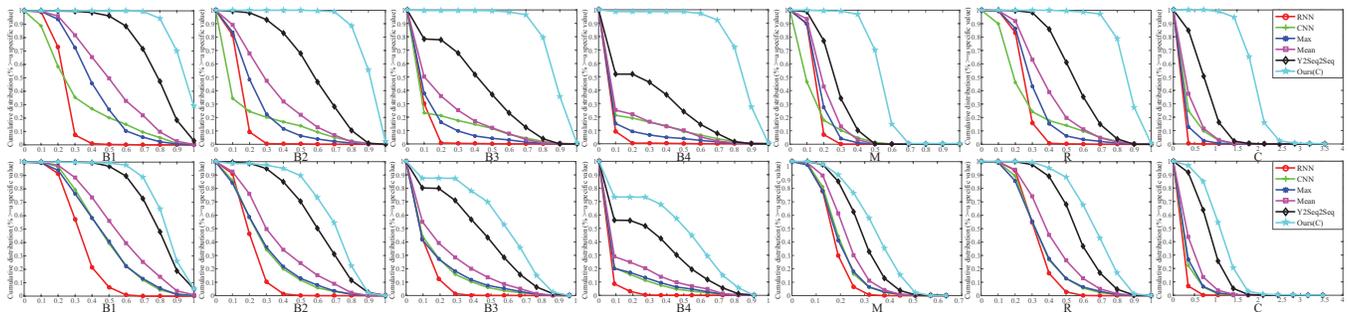


Figure 11: The comparison with methods in Table 5, Table 6 and state-of-the-art under the test set of chair (1st row) and table (2nd row) class. The comparison is conducted in terms of cumulative distribution of different metrics.

REFERENCES

- [1] Kevin Chen, Christopher B Choy, Manolis Savva, Angel X Chang, Thomas Funkhouser, and Silvio Savarese. 2018. Text2Shape: Generating Shapes from Natural Language by Learning Joint Embeddings. In *ACCV*.
- [2] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the Properties of Neural Machine Translation: Encoder-Decoder Approaches. In *SSST@EMNLP*. 103–111.
- [3] Michael Denkowski and Alon Lavie. 2014. Meteor Universal: Language Specific Translation Evaluation for Any Target Language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*. 376–380.
- [4] Zhe Gan, Chuang Gan, Xiaodong He, Yunchen Pu, Kenneth Tran, Jianfeng Gao, Lawrence Carin, and Li Deng. 2017. Semantic Compositional Networks for Visual Captioning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [5] Ross Girshick. 2015. Fast R-CNN. In *IEEE International Conference on Computer Vision*. 1440–1448.
- [6] Zhizhong Han, Chao Chen, Yu-Shen Liu, and Matthias Zwicker. 2020. DRWR: A Differentiable Renderer without Rendering for Unsupervised 3D Structure Learning from Silhouette Images. In *ICML*.
- [7] Zhizhong Han, Xinhai Liu, Yu-Shen Liu, and Matthias Zwicker. 2019. Parts4Feature: Learning 3D Global Features from Generally Semantic Parts in Multiple Views. In *IJCAI*.
- [8] Zhizhong Han, Zhenbao Liu, Junwei Han, Chi-Man Vong, Shuhui Bu, and C.L.Philip Chen. 2017. Mesh Convolutional Restricted Boltzmann Machines for Unsupervised Learning of Features With Structure Preservation on 3D Meshes. *IEEE Transactions on Neural Network and Learning Systems* 28, 10 (2017), 2268 – 2281.
- [9] Zhizhong Han, Zhenbao Liu, Junwei Han, Chi-Man Vong, Shuhui Bu, and C.L.P. Chen. 2019. Unsupervised Learning of 3D Local Features from Raw Voxels Based on A Novel Permutation Voxelization Strategy. *IEEE Transactions on Cybernetics* 49, 2 (2019), 481–494.
- [10] Zhizhong Han, Zhenbao Liu, Junwei Han, Chi-Man Vong, Shuhui Bu, and Xuelong Li. 2016. Unsupervised 3D Local Feature Learning by Circle Convolutional Restricted Boltzmann Machine. *IEEE Transactions on Image Processing* 25, 11 (2016), 5331–5344.
- [11] Zhizhong Han, Zhenbao Liu, Chi-Man Vong, Yu-Shen Liu, Shuhui Bu, Junwei Han, and C.L.Philip Chen. 2017. BoSCC: Bag of Spatial Context Correlations for Spatially Enhanced 3D Shape Representation. *IEEE Transactions on Image Processing* 26, 8 (2017), 3707–3720.
- [12] Zhizhong Han, Zhenbao Liu, Chi-Man Vong, Yu-Shen Liu, Shuhui Bu, Junwei Han, and CL Philip Chen. 2018. Deep Spatiality: Unsupervised Learning of Spatially-Enhanced Global and Local 3D Features by Deep Neural Network with Coupled Softmax. *IEEE Transactions on Image Processing* 27, 6 (2018), 3049–3063.
- [13] Zhizhong Han, Honglei Lu, Zhenbao Liu, Chi-Man Vong, Yu-Shen Liu, Matthias Zwicker, Junwei Han, and C.L. Philip Chen. 2019. 3D2SeqViews: Aggregating Sequential Views for 3D Global Feature Learning by CNN With Hierarchical Attention Aggregation. *IEEE Transactions on Image Processing* 28, 8 (2019), 3986–3999.
- [14] Zhizhong Han, Guanhuai Qiao, Yu-Shen Liu, and Matthias Zwicker. 2020. SeqXY2SeqZ: Structure Learning for 3D Shapes by Sequentially Predicting 1D Occupancy Segments From 2D Coordinates. *ArXiv abs/2003.05559* (2020).
- [15] Zhizhong Han, Mingyang Shang, Zhenbao Liu, Chi-Man Vong, Yu-Shen Liu, Matthias Zwicker, Junwei Han, and C.L. Philip Chen. 2019. SeqViews2SeqLabels: Learning 3D Global Features via Aggregating Sequential Views by RNN With Attention. *IEEE Transactions on Image Processing* 28, 2 (2019), 685–672.
- [16] Zhizhong Han, Mingyang Shang, Xiyang Wang, Yu-Shen Liu, and Matthias Zwicker. 2019. Y2Seq2Seq: Cross-Modal Representation Learning for 3D Shape and Text by Joint Reconstruction and Prediction of View and Word Sequences. In *AAAI*. 126–133.
- [17] Zhizhong Han, Xiyang Wang, Yu-Shen Liu, and Matthias Zwicker. 2019. 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 *ICCV*.
- [18] Zhizhong Han, Xiyang Wang, Chi-Man Vong, Yu-Shen Liu, Matthias Zwicker, and C.L. Philip Chen. 2019. 3DViewGraph: Learning Global Features for 3D Shapes from A Graph of Unordered Views with Attention. In *IJCAI*.
- [19] Tao Hu, Zhizhong Han, Abhinav Shrivastava, and Matthias Zwicker. 2019. Render4Completion: Synthesizing Multi-view Depth Maps for 3D Shape Completion. *ArXiv abs/1904.08366* (2019).
- [20] Tao Hu, Zhizhong Han, and Matthias Zwicker. 2020. 3D Shape Completion with Multi-view Consistent Inference. In *AAAI*.
- [21] Tao Hu, Geng Lin, Zhizhong Han, and Matthias Zwicker. 2019. Learning to Generate Dense Point Clouds with Textures on Multiple Categories. *ArXiv abs/1912.10545* (2019).
- [22] Qiuyuan Huang, Pengchuan Zhang, Dapeng Wu, and Lei Zhang. 2018. Turbo Learning for CaptionBot and DrawingBot. In *Advances in Neural Information Processing Systems*. 6455–6465.
- [23] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. 2020. SDFDiff: Differentiable Rendering of Signed Distance Fields for 3D Shape Optimization. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [24] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. 2016. DenseCap: Fully Convolutional Localization Networks for Dense Captioning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [25] Evangelos Kalogerakis, Melinos Averkiou, Subhransu Maji, and Siddhartha Chaudhuri. 2017. 3D Shape Segmentation with Projective Convolutional Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*. 6630–6639.
- [26] Andrej Karpathy and Li Fei-Fei. 2017. Deep Visual-Semantic Alignments for Generating Image Descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 4 (2017), 664–676.
- [27] Andrej Karpathy and Fei-Fei Li. 2015. Deep visual-semantic alignments for generating image descriptions. In *IEEE Conference on Computer Vision and Pattern Recognition*. 3128–3137.
- [28] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of summaries. In *Proceeding of ACL workshop on Text Summarization Branches Out*.
- [29] Xinhai Liu, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. 2019. Point2Sequence: Learning the Shape Representation of 3D Point Clouds with an Attention-based Sequence to Sequence Network. In *AAAI*. 8778–8785.
- [30] Xinhai Liu, Zhizhong Han, Wen Xin, Yu-Shen Liu, and Matthias Zwicker. 2019. L2G Auto-encoder: Understanding Point Clouds by Local-to-Global Reconstruction with Hierarchical Self-Attention. In *ACMMM*.
- [31] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy Networks: Learning 3D Reconstruction in Function Space. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [32] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Annual Meeting on Association for Computational Linguistics*. 311–318.
- [33] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [34] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. 2017. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in Neural Information Processing Systems*. 5105–5114.
- [35] Shaoqing Ren, Kaifeng He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards Real-time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems*. 91–99.
- [36] Xu Shen, Xinmei Tian, Jun Xing, Yong Rui, and Dacheng Tao. 2018. Sequence-to-Sequence Learning via Shared Latent Representation. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16071>
- [37] K. Simonyan and A. Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR abs/1409.1556* (2014).
- [38] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. 2016. Deep Metric Learning via Lifted Structured Feature Embedding. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4004–4012.
- [39] Yale Song and Mohammad Soleymani. 2018. Cross-Modal Retrieval with Implicit Concept Association. *CoRR abs/1804.04318*

- (2018). arXiv:1804.04318 <http://arxiv.org/abs/1804.04318>
- [40] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. CIDEr: Consensus-based image description evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4566–4575.
- [41] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to Sequence – Video to Text. In *Proceedings of the IEEE International Conference on Computer Vision*. 4534–4542.
- [42] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A Neural Image Caption Generator. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [43] Josiah Wang, Pranava Swaroop Madhyastha, and Lucia Specia. 2018. Object Counts! Bringing Explicit Detections Back into Image Captioning. In *NAACL-HLT*. 2180–2193.
- [44] Xin Wen, Zhizhong Han, Geunhyuk Youk, and Yu-Shen Liu. 2020. CF-SIS: Semantic-Instance Segmentation of 3D Point Clouds by Context Fusion with Self-Attention. In *ACM International Conference on Multimedia*.
- [45] Xin Wen, Tianyang Li, Zhizhong Han, and Yu-Shen Liu. 2020. Point Cloud Completion by Skip-attention Network with Hierarchical Folding. In *The IEEE Conference on Computer Vision and Pattern Recognition*.
- [46] Q. Wu, C. Shen, L. Liu, A. Dick, and A. v. d. Hengel. 2016. What Value Do Explicit High Level Concepts Have in Vision to Language Problems?. In *IEEE Conference on Computer Vision and Pattern Recognition*. 203–212.
- [47] Jin Xie, Guoxian Dai, and Yi Fang. 2017. Deep multimetric learning for shape-based 3D model retrieval. *IEEE Transactions on Multimedia* 19, 11 (2017), 2463–2474.
- [48] Xuwang Yin and Vicente Ordonez. 2017. Obj2Text: Generating Visually Descriptive Language from Object Layouts. In *EMNLP*. 177–187.
- [49] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. 2016. Image Captioning with Semantic Attention. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4651–4659.