This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Pix2Vox: Context-aware 3D Reconstruction from Single and Multi-view Images

Haozhe Xie[†] Hongxun Yao[†] Xiaoshuai Sun[†] Shangchen Zhou[‡] Shengping Zhang^{†§} [†]Harbin Institute of Technology [‡]SenseTime Research [§]Peng Cheng Laboratory

{hzxie, h.yao, xiaoshuaisun, s.zhang}@hit.edu.cn[‡]zhoushangchen@sensetime.com https://haozhexie.com/project/pix2vox

Abstract

Recovering the 3D representation of an object from single-view or multi-view RGB images by deep neural networks has attracted increasing attention in the past few years. Several mainstream works (e.g., 3D-R2N2) use recurrent neural networks (RNNs) to fuse multiple feature maps extracted from input images sequentially. However, when given the same set of input images with different orders, RNN-based approaches are unable to produce consistent reconstruction results. Moreover, due to long-term memory loss, RNNs cannot fully exploit input images to refine reconstruction results. To solve these problems, we propose a novel framework for single-view and multi-view 3D reconstruction, named Pix2Vox. By using a well-designed encoder-decoder, it generates a coarse 3D volume from each input image. Then, a context-aware fusion module is introduced to adaptively select high-quality reconstructions for each part (e.g., table legs) from different coarse 3D volumes to obtain a fused 3D volume. Finally, a refiner further refines the fused 3D volume to generate the final output. Experimental results on the ShapeNet and Pix3D benchmarks indicate that the proposed Pix2Vox outperforms state-ofthe-arts by a large margin. Furthermore, the proposed method is 24 times faster than 3D-R2N2 in terms of backward inference time. The experiments on ShapeNet unseen 3D categories have shown the superior generalization abilities of our method.

1. Introduction

3D reconstruction is an important problem in robotics, CAD, virtual reality and augmented reality. Traditional methods, such as Structure from Motion (SfM) [14] and Simultaneous Localization and Mapping (SLAM) [6], match image features across views. However, establishing feature correspondences becomes extremely difficult when multiple viewpoints are separated by a large margin due to local appearance changes or self-occlusions [12]. To overcome these limitations, several deep learning based approaches, including 3D-R2N2 [2], LSM [9], and 3DensiNet [27], have



Figure 1: Forward inference time, model size, and IoU of state-of-the-arts and our methods for single-view 3D reconstruction on the ShapeNet testing set. The radius of each circle represents the size of the corresponding model. Pix2Vox outperforms state-of-the-arts in forward inference time and reaches the best balance between accuracy and model size.

been proposed to recover the 3D shape of an object and obtained promising results.

To generate 3D volumes, 3D-R2N2 [2] and LSM [9] formulate multi-view 3D reconstruction as a sequence learning problem and use recurrent neural networks (RNNs) to fuse multiple feature maps extracted by a shared encoder from input images. The feature maps are incrementally refined when more views of an object are available. However, RNN-based methods suffer from three limitations. First, when given the same set of images with different orders, RNNs are unable to estimate the 3D shape of an object consistently results due to permutation variance [26]. Second, due to long-term memory loss of RNNs, the input images cannot be fully exploited to refine reconstruction results [15]. Last but not least, RNN-based methods are timeconsuming since input images are processed sequentially



Figure 2: An overview of the proposed Pix2Vox. The network recovers the shape of 3D objects from arbitrary (uncalibrated) single or multiple images. The reconstruction results can be refined when more input images are available. Note that the weights of the encoder and decoder are shared among all views.

without parallelization [8].

To address the issues mentioned above, we propose Pix2Vox, a novel framework for single-view and multi-view 3D reconstruction that contains four modules: encoder, decoder, context-aware fusion, and refiner. The encoder and decoder generate coarse 3D volumes from multiple input images in parallel, which eliminates the effect of the orders of input images and accelerates the computation. Then, the context-aware fusion module selects high-quality reconstructions from all coarse 3D volumes and generates a fused 3D volume, which fully exploits information of all input images without long-term memory loss. Finally, the refiner further correct wrongly recovered parts of the fused 3D volumes to obtain a refined reconstruction. To achieve a good balance between accuracy and model size, we implement two versions of the proposed framework: Pix2Vox-F and Pix2Vox-A (Figure 1).

The contributions can be summarized as follows:

- We present a unified framework for both single-view and multi-view 3D reconstruction, namely Pix2Vox. We equip Pix2Vox with well-designed encoder, decoder, and refiner, which shows a powerful ability to handle 3D reconstruction in both synthetic and realworld images.
- We propose a context-aware fusion module to adaptively select high-quality reconstructions for each part from different coarse 3D volumes in parallel to produce a fused reconstruction of the whole object. To the best of our knowledge, it is the first time to exploit context across multiple views for 3D reconstruction.
- Experimental results on the ShapeNet [33] and Pix3D [22] datasets demonstrate that the proposed approaches outperform state-of-the-art methods in terms of both accuracy and efficiency. Additional experiments also show its strong generalization abilities in reconstructing unseen 3D objects.

2. Related Work

Single-view 3D Reconstruction Theoretically, recovering 3D shape from single-view images is an ill-posed problem. To address this issue, many attempts have been made, such as ShapeFromX [1, 18], where X may represent silhouettes [4], shading [16], and texture [30]. However, these methods are barely applicable to use in the real-world scenarios, because all of them require strong presumptions and abundant expertise in natural images [35]. With the success of generative adversarial networks (GANs) [7] and variational autoencoders (VAEs) [11], 3D-VAE-GAN [32] adopts GAN and VAE to generate 3D objects by taking a single-view image as input. However, 3D-VAE-GAN requires class labels for reconstruction. MarrNet [31] reconstructs 3D objects by estimating depth, surface normals, and silhouettes of 2D images, which is challenging and usually leads to severe distortion [24]. OGN [23] and O-CNN [29] use octree to represent higher resolution volumetric 3D objects with a limited memory budget. However, OGN representations are complex and consume more computational resources due to the complexity of octree representations. PSGN [5] and 3D-LMNet [13] generate point clouds from single-view images. However, the points have a large degree of freedom in the point cloud representation because of the limited connections between points. Consequently, these methods cannot recover 3D volumes accurately [28].

Multi-view 3D Reconstruction SfM [14] and SLAM [6] methods are successful in handling many scenarios. These methods match features among images and estimate the camera pose for each image. However, the matching process becomes difficult when multiple viewpoints are separated by a large margin. Besides, scanning all surfaces of an object before reconstruction is sometimes impossible, which leads to incomplete 3D shapes with occluded or hollowed-out areas [34]. Powered by large-scale datasets of 3D CAD models (e.g., ShapeNet [33]), deep-learning-



Figure 3: The network architecture of (top) Pix2Vox-F and (bottom) Pix2Vox-A. The EDLoss and the RLoss are defined as Equation 3. To reduce the model size, the refiner is removed in Pix2Vox-F.

based methods have been proposed for 3D reconstruction. Both 3D-R2N2 [2] and LSM [9] use RNNs to infer 3D shape from single or multiple input images and achieve impressive results. However, RNNs are time-consuming and permutation-variant, which produce inconsistent reconstruction results. 3DensiNet [27] uses max pooling to aggregate the features from multiple images. However, max pooling only extracts maximum values from features, which may ignore other valuable features that are useful for 3D reconstruction.

3. The Method

3.1. Overview

The proposed Pix2Vox aims to reconstruct the 3D shape of an object from either single or multiple RGB images. The 3D shape of an object is represented by a 3D voxel grid, where 0 is an empty cell and 1 denotes an occupied cell. The key components of Pix2Vox are shown in Figure 2. First, the encoder produces feature maps from input images. Second, the decoder takes each feature map as input and generates a coarse 3D volume correspondingly. Third, single or multiple 3D volumes are forwarded to the contextaware fusion module, which adaptively selects high-quality reconstructions for each part from coarse 3D volumes to obtain a fused 3D volume. Finally, the refiner with skipconnections further refines the fused 3D volume to generate the final reconstruction result.

3.2. Network Architecture

Figure 3 shows the detailed architectures of Pix2Vox-F and Pix2Vox-A. The former involves much fewer parameters and lower computational complexity, while the latter has more parameters, which can construct more accurate 3D shapes but has higher computational complexity.

3.2.1 Encoder

The encoder is to compute a set of features for the decoder to recover the 3D shape of the object. The first nine convolutional layers, along with the corresponding batch normalization layers and ReLU activations of a VGG16 [20] pretrained on ImageNet [3], are used to extract a $512 \times 28 \times 28$ feature tensor from a $224 \times 224 \times 3$ image. This feature extraction is followed by three sets of 2D convolutional layers, batch normalization layers and ELU layers to embed semantic information into feature vectors. In Pix2Vox-F, the kernel size of the first convolutional layer is 1^2 while the kernel sizes of the other two are 3^2 . The number of output channels of the convolutional layer starts with 512 and decreases by half for the subsequent layer and ends up with 128. In Pix2Vox-A, the kernel sizes of the three convolutional layers are 3^2 , 3^2 , and 1^2 , respectively. The output channels of the three convolutional layers are 512, 512, and 256, respectively. After the second convolutional layer, there is a max pooling layer with kernel sizes of 3^2 and 4^2 in Pix2Vox-F and Pix2Vox-A, respectively. The feature vectors produced by Pix2Vox-F and Pix2Vox-A are of sizes 2048 and 16384, respectively.

3.2.2 Decoder

The decoder is responsible for transforming information of 2D feature maps into 3D volumes. There are five 3D transposed convolutional layers in both Pix2Vox-F and Pix2Vox-A. Specifically, the first four transposed convolutional layers are of a kernel size of 4^3 , with stride of 2 and padding of 1. There is an additional transposed convolutional layer with a bank of 1^3 filter. Each transposed convolutional layer is followed by a batch normalization layer and a ReLU activation except for the last layer followed by a sigmoid function. In Pix2Vox-F, the numbers of output channels of the transposed convolutional layers are 128, 64, 32, 8, and 1, re-



Figure 4: Visualization of the score maps in the contextaware fusion module. The context-aware fusion module generates higher scores for high-quality reconstructions, which can eliminate the effect of the missing or wrongly recovered parts.



Figure 5: An overview of the context-aware fusion module. It aims to select high-quality reconstructions for each part to construct the final results. The objects in the bounding box describe the procedure score calculation for a coarse volume v_n^c . The other scores are calculated according to the same procedure. Note that the weights of the context scoring network are shared among different views.

spectively. In Pix2Vox-A, the numbers of output channels of the five transposed convolutional layers are 512, 128, 32, 8, and 1, respectively. The decoder outputs a 32^3 voxelized shape in the object's canonical view.

3.2.3 Context-aware Fusion

From different viewpoints, we can see different visible parts of an object. The reconstruction qualities of visible parts are much higher than those of invisible parts. Inspired by this observation, we propose a context-aware fusion module to adaptively select high-quality reconstruction for each part (e.g., table legs) from different coarse 3D volumes. The selected reconstructions are fused to generate a 3D volume of the whole object (Figure 4).

As shown in Figure 5, given coarse 3D volumes and the corresponding context, the context-aware fusion module generates a score map for each coarse volume and then fuses them into one volume by the weighted summation of all coarse volumes according to their score maps. The spatial information of voxels is preserved in the context-aware fusion module, and thus Pix2Vox can utilize multi-view information to recover the structure of an object better.

Specifically, the context-aware fusion module generates the context c_r of the r-th coarse volume v_r^c by concatenating the output of the last two layers in the decoder. Then, the context scoring network generates a score m_r for the context of the r-th coarse voxel. The context scoring network is composed of five sets of 3D convolutional layers, each of which has a kernel size of 3^3 and padding of 1, followed by a batch normalization and a leaky ReLU activation. The numbers of output channels of convolutional layers are 9, 16, 8, 4, and 1, respectively. The learned score m_r for context c_r are normalized across all learnt scores. We choose softmax as the normalization function. Therefore, the score $s_r^{(i,j,k)}$ at position (i, j, k) for the r-th voxel can be calculated as

$$s_{r}^{(i,j,k)} = \frac{\exp\left(m_{r}^{(i,j,k)}\right)}{\sum_{p=1}^{n} \exp\left(m_{p}^{(i,j,k)}\right)}$$
(1)

where n represents the number of views. Finally, the fused voxel v^f is produced by summing up the product of coarse voxels and the corresponding scores altogether.

$$v^f = \sum_{r=1}^n s_r v_r^c \tag{2}$$

3.2.4 Refiner

The refiner can be seen as a residual network, which aims to correct wrongly recovered parts of a 3D volume. It follows the idea of a 3D encoder-decoder with the U-net connections [17]. With the help of the U-net connections between the encoder and decoder, the local structure in the fused volume can be preserved. Specifically, the encoder has three 3D convolutional layers, each of which has a bank of 4^3 filters with padding of 2, followed by a batch normalization layer, a leaky ReLU activation and a max pooling layer with a kernel size of 2^3 . The numbers of output channels of convolutional layers are 32, 64, and 128, respectively. The encoder is finally followed by two fully connected layers with dimensions of 2048 and 8192. The decoder consists of three transposed convolutional layers, each of which has a bank of 4^3 filters with padding of 2 and stride of 1.

Except for the last transposed convolutional layer that is followed by a sigmoid function, other layers are followed by a batch normalization layer and a ReLU activation.

Table 1: Single-view reconstruction on ShapeNet compared using Intersection-over-Union (IoU). The best number for each category is highlighted in bold. Note that DRC [25] is trained/tested per category and PSGN [5] takes object masks as an additional input. Besides, PSGN uses 220k 3D CAD models while the remaining methods use only 44k 3D CAD models during training.

Category	3D-R2N2 [2]	OGN [23]	DRC [25]	PSGN [5]	Pix2Vox-F	Pix2Vox-A
airplane	0.513	0.587	0.571	0.601	0.600	0.684
bench	0.421	0.481	0.453	0.550	0.538	0.616
cabinet	0.716	0.729	0.635	0.771	0.765	0.792
car	0.798	0.828	0.755	0.831	0.837	0.854
chair	0.466	0.483	0.469	0.544	0.535	0.567
display	0.468	0.502	0.419	0.552	0.511	0.537
lamp	0.381	0.398	0.415	0.462	0.435	0.443
speaker	0.662	0.637	0.609	0.737	0.707	0.714
rifle	0.544	0.593	0.608	0.604	0.598	0.615
sofa	0.628	0.646	0.606	0.708	0.687	0.709
table	0.513	0.536	0.424	0.606	0.587	0.601
telephone	0.661	0.702	0.413	0.749	0.770	0.776
watercraft	0.513	0.632	0.556	0.611	0.582	0.594
Overall	0.560	0.596	0.545	0.640	0.634	0.661

Table 2: Multi-view reconstruction on ShapeNet compared using Intersection-over-Union (IoU). The best results for different numbers of views are highlighted in bold. The marker † indicates that the context-aware fusion is replaced with the average fusion.

Methods	1 view	2 views	3 views	4 views	5 views	8 views	12 views	16 views	20 views
3D-R2N2 [2]	0.560	0.603	0.617	0.625	0.634	0.635	0.636	0.636	0.636
Pix2Vox-F [†]	0.634	0.653	0.661	0.666	0.668	0.672	0.674	0.675	0.676
Pix2Vox-F	0.634	0.660	0.668	0.673	0.676	0.680	0.682	0.684	0.684
Pix2Vox-A [†]	0.661	0.678	0.684	0.687	0.689	0.692	0.694	0.695	0.695
Pix2Vox-A	0.661	0.686	0.693	0.697	0.699	0.702	0.704	0.705	0.706

3.2.5 Loss Function

The loss function of the network is defined as the mean value of the voxel-wise binary cross entropies between the reconstructed object and the ground truth. More formally, it can be defined as

$$\ell = \frac{1}{N} \sum_{i=1}^{N} \left[gt_i \log(p_i) + (1 - gt_i) \log(1 - p_i) \right]$$
(3)

where N denotes the number of voxels in the ground truth. p_i and gt_i represent the predicted occupancy and the corresponding ground truth. The smaller the ℓ value is, the closer the prediction is to the ground truth.

4. Experiments

4.1. Datasets and Metrics

Datasets We evaluate the proposed Pix2Vox-F and Pix2Vox-A on both synthetic images of objects from the

ShapeNet [33] dataset and real images from the Pix3D [22] dataset. More specifically, we use a subset of ShapeNet consisting of 13 major categories and 43,783 3D models following the settings of [2]. As for Pix3D, we use the 2,894 untruncated and unoccluded chair images following the settings of [22].

Evaluation Metrics To evaluate the quality of the output from the proposed methods, we binarize the probabilities at a fixed threshold of 0.3 and use intersection over union (IoU) as the similarity measure. More formally,

$$IoU = \frac{\sum_{i,j,k} I(p_{(i,j,k)} > t) I(gt_{(i,j,k)})}{\sum_{i,j,k} I \left[I(p_{(i,j,k)} > t) + I(gt_{(i,j,k)}) \right]}$$
(4)

where $p_{(i,j,k)}$ and $gt_{(i,j,k)}$ represent the predicted occupancy probability and the ground truth at (i, j, k), respectively. $I(\cdot)$ is an indicator function and t denotes a voxelization threshold. Higher IoU values indicate better reconstruction results.



Figure 6: Single-view (left) and multi-view (right) reconstructions on the ShapeNet testing set. GT represents the ground truth of the 3D object. Note that DRC [25] is trained/tested per category.

4.2. Implementation Details

We use 224×224 RGB images as input to train the proposed methods with a shape batch size of 64. The output voxelized reconstruction is 32^3 in size. We implement our network in PyTorch and train both Pix2Vox-F and Pix2Vox-A using an Adam optimizer [10] with a β_1 of 0.9 and a β_2 of 0.999. The initial learning rate is set to 0.001 and decayed by 2 after 150 epochs. First, we train both networks except the context-aware fusion feeding with a single-view image for 250 epochs. Then, we train the whole network jointly feeding with random numbers of input images for 100 epochs.

4.3. Reconstruction of Synthetic Images

To evaluate the performance of the proposed methods in handling synthetic images, we compare our methods against several state-of-the-art methods on the ShapeNet testing set. To make a fair comparison, all methods are compared with the same input images for all experiments except PSGN [5]. Although PSGN uses much more data during training, Pix2Vox-A still performs better in recovering the 3D shape of an object. Table 1 shows the performance of single-view reconstruction, while Table 2 shows the mean IoU scores of multi-view reconstruction with different numbers of views.

The single-view reconstruction results of Pix2Vox-F and Pix2Vox-A significantly outperform other methods (Table 1). Pix2Vox-A increases IoU over 3D-R2N2 by 18%. In multi-view reconstruction, Pix2Vox-A consistently outperforms 3D-R2N2 in all numbers of views (Table 2). The IoU of Pix2Vox-A is 13% higher than that of 3D-R2N2.

Table 3: Single-view reconstruction on Pix3D compared using Intersection-over-Union (IoU). The best number is highlighted in bold.

Method	IoU
3D-R2N2 [2]	0.136
DRC [25]	0.265
Pix3D (w/o Pose) [22]	0.267
Pix3D (w/ Pose) [22]	0.282
Pix2Vox-F	0.271
Pix2Vox-A	0.288

Figure 6 shows several reconstruction examples from the ShapeNet testing set. Both Pix2Vox-F and Pix2Vox-A are able to recover the thin parts of objects, such as lamps and table legs. Compare with Pix2Vox-F, we also observe that higher dimensional feature maps in Pix2Vox-A do contribute to 3D reconstruction. Moreover, in multi-view reconstruction, both Pix2Vox-A and Pix2Vox-F produce better results than 3D-R2N2.

4.4. Reconstruction of Real-world Images

To evaluate the performance on of the proposed methods on real-world images, we test our methods for single-view reconstruction on the Pix3D dataset.

We use the pipeline of RenderForCNN [21] to generate 60 images for each 3D CAD model in the ShapeNet dataset. We perform quantitative evaluation of the resulting models on real-world RGB images using the Pix3D dataset. Besides, we augment our training data by random color and



Figure 7: Reconstruction on the Pix3D testing set from single-view images. GT represents the ground truth of the 3D object.

light jittering. First, the images are cropped according to the bounding box of the objects within the image. Then, these cropped images are rescaled as required by each reconstruction network.

The mean IoU of the Pix3D dataset is reported in Table 3. The experimental results indicate Pix2Vox-A outperform the competing approaches on the Pix3D testing set without estimating the pose of an object. The qualitative analysis is given in Figure 7, which indicate that the proposed methods are more effective in handling real-world scenarios.

4.5. Reconstruction of Unseen Objects

In order to test how well our methods can generalize to unseen objects, we conduct additional experiments on ShapeNetCore [33]. We use Mitsuba¹ to render objects in the remaining 44 categories of ShapeNetCore from 24 random views along with voxel representations. All pretrained models have never "seen" either the objects in these categories or the labels of objects before. More specifically, all models are trained on the 13 major categories of ShapeNet renderings provided by [2] and tested on the remaining 44 categories of ShapeNetCore with the same input images. The reconstruction results of 3D-R2N2 are obtained with the released pretrained model.

Several reconstruction results are presented in Figure 8. The reconstruction IoU of 3D-R2N2 on unseen objects



Figure 8: Reconstruction on unseen objects of ShapeNet from 5-view images. GT represents the ground truth of the 3D object.

is 0.120, while Pix2Vox-F and Pix2Vox-A are 0.209 and 0.227, respectively. Experimental results demonstrate that 3D-R2N2 can hardly recover the shape of unseen objects. In contrast, Pix2Vox-F and Pix2Vox-A show satisfactory generalization abilities to unseen objects.

4.6. Ablation Study

In this section, we validate the context-aware fusion and the refiner by ablation studies.

Context-aware fusion To quantitatively evaluate the context-aware fusion, we replace the context-aware fusion in Pix2Vox-A with the average fusion, where the fused voxel v^f can be calculated as

$$v_{(i,j,k)}^{f} = \frac{1}{n} \sum_{r=1}^{n} v_{(i,j,k)}^{r}$$
(5)

Table 2 shows that the context-aware fusion performs better than the average fusion in selecting the high-quality reconstructions for each part from different coarse volumes.

To make a further comparison with RNN-based fusion, we remove the context-aware fusion and add an 3D convolutional LSTM [2] after the encoder. To fit the input of the 3D convolutional LSTM, we add an additional fully-connected layer with a dimension of 1024 before it. As shown in Figure 9a, both the average fusion and context-aware fusion consistently outperform the RNN-based fusion in all numbers of views.

¹https://www.mitsuba-renderer.org



Figure 9: The IoU on ShapeNet testing set. (a) Effects of the context aware fusion and the number of views on the evaluation IoU. (b) Effects of the refiner network and the number of views on the evaluation IoU.

Table 4: Memory usage and running time on ShapeNet dataset. Note that backward time is measured in single-view reconstruction with a batch size of 1.

Methods	3D-R2N2	OGN	Pix2Vox-F	Pix2Vox-A
#Parameters (M)	35.97	12.46	7.41	114.24
Memory (MB)	1407	793	673	2729
Training (hours)	169	192	12	25
Backward (ms)	312.50	312.25	12.93	72.01
Forward, 1-view (ms)	73.35	37.90	9.25	9.90
Forward, 2-views (ms)	108.11	N/A	12.05	13.69
Forward, 4-views (ms)	112.36	N/A	23.26	26.31
Forward, 8-views (ms)	117.64	N/A	52.63	55.56

Refiner Pix2Vox-A uses a refiner to further refine the fused 3D volume. For single-view reconstruction on ShapeNet, the IoU of Pix2Vox-A is 0.661. In contrast, the IoU of Pix2Vox-A without the refiner decreases to 0.636. Removing refiner causes considerable degeneration for the reconstruction accuracy. As shown in Figure 9b, as the number of views increases, the effect of the refiner becomes weaker.

The ablation studies indicate that both the context-aware fusion and the refiner play important roles in our framework for the performance improvements against previous stateof-the-art methods.

4.7. Space and Time Complexity

Table 4 and Figure 1 show the numbers of parameters of different methods. There is an 80% reduction in parameters in Pix2Vox-F compared to 3D-R2N2.

The running times are obtained on the same PC with an NVIDIA GTX 1080 Ti GPU. For more precise timing, we exclude the reading and writing time when evaluating the forward and backward inference time. Both Pix2Vox-F and Pix2Vox-A are about 8 times faster in forward inference than 3D-R2N2 in single-view reconstruction. In backward inference, Pix2Vox-F and Pix2Vox-A are about 24 and 4 times faster than 3D-R2N2, respectively.

4.8. Discussion

To give a detailed analysis of the context-aware fusion module, we visualized the score maps of three coarse volumes when reconstructing the 3D shape of a table from 3view images, as shown in Figure 4. The reconstruction quality of the table tops on the right is clearly of low quality, and the score of the corresponding part is lower than those in the other two coarse volumes. The fused 3D volume is obtained by combining the selected high-quality reconstruction parts, where bad reconstructions can be eliminated effectively by our scoring scheme.

Pix2Vox recovers the 3D shape of an object without knowing camera parameters. To further demonstrate the superior ability of the context-aware fusion in multi-view stereo (MVS) systems [19], we replace the RNN with the context-aware fusion in LSM [9]. Specifically, we remove the recurrent fusion and add the context-aware fusion to combine the 3D volume reconstruction of each view. Experimental results show that the IoU is increased by about 2% on the ShapeNet testing set, which indicate that the contextaware fusion also helps MVS systems to obtain better reconstruction results.

Although our methods outperform state-of-the-arts, the reconstruction results of our methods are still with a low resolution. We can further improve the reconstruction resolutions in the future work by introducing GANs [7].

5. Conclusion and Future Works

In this paper, we propose a unified framework for both single-view and multi-view 3D reconstruction, named Pix2Vox. Compared with existing methods that fuse deep features generated by a shared encoder, the proposed method fuses multiple coarse volumes produced by a decoder and preserves multi-view spatial constraints better. Quantitative and qualitative evaluation for both single-view and multi-view reconstruction on the ShapeNet and Pix3D benchmarks indicate that the proposed methods outperform state-of-the-arts by a large margin. Pix2Vox is computationally efficient, which is 24 times faster than 3D-R2N2 in terms of backward inference time. In future work, we will work on improving the resolution of the reconstructed 3D objects. In addition, we also plan to extend Pix2Vox to reconstruct 3D objects from RGB-D images.

Acknowledgements This work was supported by the National Natural Science Foundation of China under Project No. 61772158, 61702136, 61872112 and U1711265. We gratefully acknowledge Prof. Junbao Li and Huanyu Liu for providing additional GPU hours for this research. We would also like to thank Prof. Wangmeng Zuo, Jiapeng Tang, Xiusheng Lu, and anonymous reviewers for their valuable feedbacks and help during this research.

References

- Jonathan T. Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *TPAMI*, 37(8):1670–1687, 2015.
- [2] Christopher Bongsoo Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3D-R2N2: A unified approach for single and multi-view 3D object reconstruction. In ECCV 2016.
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In CVPR 2009.
- [4] Endri Dibra, Himanshu Jain, A. Cengiz Öztireli, Remo Ziegler, and Markus H. Gross. Human shape from silhouettes using generative HKS descriptors and cross-modal neural networks. In *CVPR 2017*.
- [5] Haoqiang Fan, Hao Su, and Leonidas J. Guibas. A point set generation network for 3D object reconstruction from a single image. In *CVPR 2017*.
- [6] Jorge Fuentes-Pacheco, José Ruíz Ascencio, and Juan Manuel Rendón-Mancha. Visual simultaneous localization and mapping: a survey. *Artif. Intell. Rev.*, 43(1):55–81, 2015.
- [7] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS* 2014.
- [8] Kyuyeon Hwang and Wonyong Sung. Single stream parallelization of generalized LSTM-like RNNs on a GPU. In *ICASSP 2015*.
- [9] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *NIPS 2017*.
- [10] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR 2015*.
- [11] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. arXiv, abs/1312.6114, 2013.
- [12] David G. Lowe. Distinctive image features from scaleinvariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [13] Priyanka Mandikal, Navaneet K. L., Mayank Agarwal, and Venkatesh Babu Radhakrishnan. 3D-LMNet: Latent embedding matching for accurate and diverse 3D point cloud reconstruction from a single image. In *BMVC 2018*.
- [14] Onur Özyeşil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion. *Acta Numerica*, 26:305–364.
- [15] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *ICML* 2013.
- [16] Stephan R. Richter and Stefan Roth. Discriminative shape from shading in uncalibrated illumination. In *CVPR 2015*.
- [17] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI 2015*.
- [18] Silvio Savarese, Marco Andreetto, Holly E. Rushmeier, Fausto Bernardini, and Pietro Perona. 3D reconstruction by shadow carving: Theory and practical evaluation. *IJCV*, 71(3):305–336, 2007.

- [19] Steven M. Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In CVPR 2006.
- [20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR* 2015.
- [21] Hao Su, Charles Ruizhongtai Qi, Yangyan Li, and Leonidas J. Guibas. Render for CNN: viewpoint estimation in images using cnns trained with rendered 3d model views. In *ICCV 2015*.
- [22] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B. Tenenbaum, and William T. Freeman. Pix3d: Dataset and methods for single-image 3d shape modeling. In *CVPR 2018*.
- [23] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3D outputs. In *ICCV 2017*.
- [24] Shubham Tulsiani. Learning Single-view 3D Reconstruction of Objects and Scenes. PhD thesis, UC Berkeley, 2018.
- [25] Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In CVPR 2017.
- [26] Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. Order matters: Sequence to sequence for sets. In *ICLR 2016*.
- [27] Meng Wang, Lingjing Wang, and Yi Fang. 3DensiNet: A robust neural network architecture towards 3D volumetric object prediction from 2D image. In ACM MM 2017.
- [28] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2Mesh: Generating 3D mesh models from single RGB images. In ECCV 2018.
- [29] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-CNN: octree-based convolutional neural networks for 3D shape analysis. ACM Trans. Graph., 36(4):72:1–72:11, 2017.
- [30] Andrew P. Witkin. Recovering surface shape and orientation from texture. *Artif. Intell.*, 17(1-3):17–45, 1981.
- [31] Jiajun Wu, Yifan Wang, Tianfan Xue, Xingyuan Sun, Bill Freeman, and Josh Tenenbaum. MarrNet: 3D shape reconstruction via 2.5D sketches. In *NIPS 2017*.
- [32] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *NIPS* 2016.
- [33] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A deep representation for volumetric shapes. In *CVPR 2015*.
- [34] Bo Yang, Stefano Rosa, Andrew Markham, Niki Trigoni, and Hongkai Wen. Dense 3D object reconstruction from a single depth view. *TPAMI*, DOI: 10.1109/TPAMI.2018.2868195, 2018.
- [35] Yang Zhang, Zhen Liu, Tianpeng Liu, Bo Peng, and Xiang Li. RealPoint3D: An efficient generation network for 3d object reconstruction from a single image. *IEEE Access*, 7:57539–57549, 2019.