

# PBP-Net: Point Projection and Back-Projection Network for 3D Point Cloud Segmentation

JuYoung Yang\*, Chanho Lee\*, Pyunghwan Ahn\*, Haeil Lee\*, Eojindl Yi\* and Junmo Kim\*

\*School of Electrical Engineering

Korea Advanced Institute of Science and Technology, Korea

Email: (yjy6711, yiwan99, dksvudghks, haeil.lee, djwld93, junmo.kim)@kaist.ac.kr

**Abstract**—Following considerable development in 3D scanning technologies, many studies have recently been proposed with various approaches for 3D vision tasks, including some methods that utilize 2D convolutional neural networks (CNNs). However, even though 2D CNNs have achieved high performance in many 2D vision tasks, existing works have not effectively applied them onto 3D vision tasks. In particular, segmentation has not been well studied because of the difficulty of dense prediction for each point, which requires rich feature representation. In this paper, we propose a simple and efficient architecture named point projection and back-projection network (PBP-Net), which leverages 2D CNNs for the 3D point cloud segmentation. 3 modules are introduced, each of which projects 3D point cloud onto 2D planes, extracts features using a 2D CNN backbone, and back-projects features onto the original 3D point cloud. To demonstrate effective 3D feature extraction using 2D CNN, we perform various experiments including comparison to recent methods. We analyze the proposed modules through ablation studies and perform experiments on object part segmentation (ShapeNet-Part dataset) and indoor scene semantic segmentation (S3DIS dataset). The experimental results show that proposed PBP-Net achieves comparable performance to existing state-of-the-art methods.

## I. INTRODUCTION

Deep learning has shown successful results in various computer vision applications such as image classification, detection, and segmentation, owing to the development in 2D convolutional neural networks (CNNs) [9], [16]. 2D CNNs can efficiently capture global and spatial features with good generalization capability. Recently, improvement in 3D scanning devices have enabled the collection of 3D data. Thus, these devices have attracted significant attention for applications in several deep learning tasks such as 3D classification, object detection, and semantic segmentation. Previously, 3D data has been represented as multi-view images [28], volumes [14], [21], and structural data [8]. Compared to multi-view images and volumetric representations, point cloud, which refers to an unordered set of points, precisely represents data and requires lesser memory. Hence, a majority of recently proposed studies attempt to effectively deal with 3D point clouds.

PointNet [20] is a pioneering research for the classification and semantic segmentation of 3D point cloud. PointNet calculates the features of individual points via symmetric multi-layer perceptrons (MLP) to guarantee equivariance to the order of points. Previous studies have employed graph convolution, which is one of the most frequently

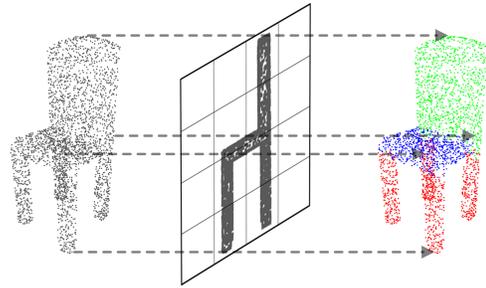


Fig. 1. Brief concept of our idea. We project 3D points onto 2D planes to extract spatial features, and re-allocate these acquired features to the original points. We actually perform part segmentation for the chair and show the result

used kernels to generate point-wise features, independent of the point order [10], [30]. Although these symmetric neural networks are capable of extracting local features of a point cloud, they often face difficulties when capturing the spatial features of point clouds. The most intuitive method to generate spatial features is directly using a 3D CNN through voxelization [17], [19]. Even though 3D CNNs are capable of capturing the spatial features of high quality, these methods require significant computational capacities, which limits the size of the voxel and its scalability. Recent studies have demonstrated improved performance over these baseline methods; however, designing specialized neural networks for 3D point clouds is still a challenging task as compared to designing 2D CNNs.

As 2D CNNs exhibit good performances in various computer vision tasks, it is beneficial to utilize 2D CNNs for 3D recognition tasks. Even though several previous studies have applied 2D convolution for the classification [25] and detection [11] from 3D point clouds, segmentation is more challenging task and requires dense predictions for each point. Also, most of their methods mainly focused on extracting global features or reasoning on 2D planes. On the other hand, we aim to extract discriminative features for each point. Some studies like [34] utilize 2D CNNs by using voxel representations which requires huge memory allocation and limit its capacity. Since 2D CNNs have an effective performance in terms of memory and computation, we maximize that advantage by simplifying each process.

In this paper, we propose a general framework of 2D CNN-based 3D point-wise feature extraction. The proposed

method is computationally efficient and easy to implement. Our core idea is to first project 3D points and their local features onto 2D planes, and subsequently extracting spatial features by using 2D convolutional layers. Finally, these spatial features are back-projected onto 3D points. Figure 1 briefly shows the concept of our idea. Based on this, we construct an efficient neural network termed as point projection and back-projection network (PBP-Net), which performs 3D point cloud segmentation tasks using 2D CNNs. Considering that 3D-to-2D projection and back-projection are differentiable, our PBP-Net is end-to-end trainable. Furthermore, PBP-Net achieves performance comparable to that of state-of-the-art methods and is applicable to various scales of datasets, including datasets with a large number of points or volumes, without any dramatic change in the network architecture.

The main contributions of this study are as follows:

- 1) We suggest 3D-to-2D projection and 2D-to-3D back-projection modules to employ 2D CNNs as 3D point-wise feature extractors. Both modules are completely differentiable; hence, we can train the overall architecture end-to-end.
- 2) PBP-Net is a simple yet powerful 2D convolution-based architecture for point cloud segmentation. Various 2D CNNs that are successful for 2D segmentation tasks can be directly utilized without any modifications. Thus, PBP-Net ensures structural freedom when choosing 2D convolutional backbones.
- 3) The performance of PBP-Net is comparable to state-of-the-art methods, in terms of part segmentation and semantic segmentation. Our experimental results demonstrate that PBP-Net is capable of extracting 3D feature representations using 2D convolution.

## II. RELATED WORKS

In this section, we discuss the previous studies related to our method. As the proposed method aims to perform semantic segmentation of point clouds by using 2D CNNs, our work is related to studies pertaining to point cloud recognition as well as those pertaining to image segmentation. First, we introduce the deep learning-based methods for irregular 3D point cloud recognition. Subsequently, we review the convolution-based methods on regular grids in 2D (images) or 3D (shapes) spaces.

**Neural Networks on irregular point clouds.** PointNet [20], which is a pioneering method for 3D point cloud recognition, acts as a simple baseline for various tasks such as classification and segmentation. PointNet employs MLP that extracts features from individual points. Thereafter, it uses the max-pooling features to produce global feature vectors. The global features can be used for object classification or they can be combined with point-wise features for segmentation (i.e., point-wise classification). PointNet achieves good performance for 3D object classification and part segmentation, and requires low computational capacities. However, follow-up studies have reported that PointNet has difficulties in capturing local features, which results from the learning interactions between neighboring points.

To extract better local features, previous works have attempted to generate local-to-global hierarchical features [6], [22] or to train the relationships between neighboring points, based on graph structures [10], [30], [32]. PointNet++ [22] was suggested for local feature extraction by applying PointNet to a set of points in the local neighborhood. Pointwise CNN [6] has been proposed to voxelize the neighborhood of a point for the application of 3D convolution on the continuous space.

Previous studies have also attempted to apply graph neural networks for point cloud recognition. [30] proposed EdgeConv operator that effectively maps the relation between the neighboring features into edge weights such that the graph convolution can infer these relationships. SpiderCNN [32] proposed a convolution filter as a product of the step function and a Taylor polynomial to capture local geodesic information while maintaining the expressiveness of the kernel. In [10], the point cloud is initially over-segmented into simple parts of the objects, which are known as superpoints; subsequently, the relationships between these groups are trained via gated recurrent unit (GRU) structures.

**Neural Networks on regular grids.** Convolutional neural networks can be used to perform 3D point cloud tasks in 2D or 3D spaces. To apply 3D convolution on point clouds, the point clouds must first be voxelized. In [17], voxels are generated as an occupancy grid with a binary value that indicates whether a voxel includes a point or not. However, 3D convolution requires significant computational resources, which prevents the voxelization of 3D spaces with high resolution. This limits the expressiveness of the voxelized representation of the point clouds, thereby limiting the performance. To overcome these problems, [23] proposed the application of PointNet to a local neighborhood prior to using the 3D convolution layers. This alleviates information loss due to abstraction that arises when voxelizing the 3D space.

Previous works have also aimed to transfer the 3D point cloud to the 2D domain and perform tasks using conventional 2D CNNs. [28] proposed the use of a 2D CNN for the multi-view images of 3D shape models. This method extracts features from 2D images and then combines them using the view-pooling layer. [25] utilized a depth map to project point clouds onto a 2D plane and then applied a CNN to perform classification on the depth image. For 3D object detection, [11] used pillar features, which were produced by applying PointNet at each pillar, as intermediate image features which were then processed via 2D convolution to produce the 3D bounding boxes.

Similarly, we also leverage conventional 2D image segmentation methods in our study. The fully convolutional network (FCN) [16] is one of the most widely used baseline for image segmentation. In the FCN, the last fully connected layers of the classification networks are replaced by (de)convolution layers such that pixel-wise predictions are produced. Following the principle of the FCN, [18] proposed a symmetric network structure wherein the first stage hierarchically encoded the features and the second

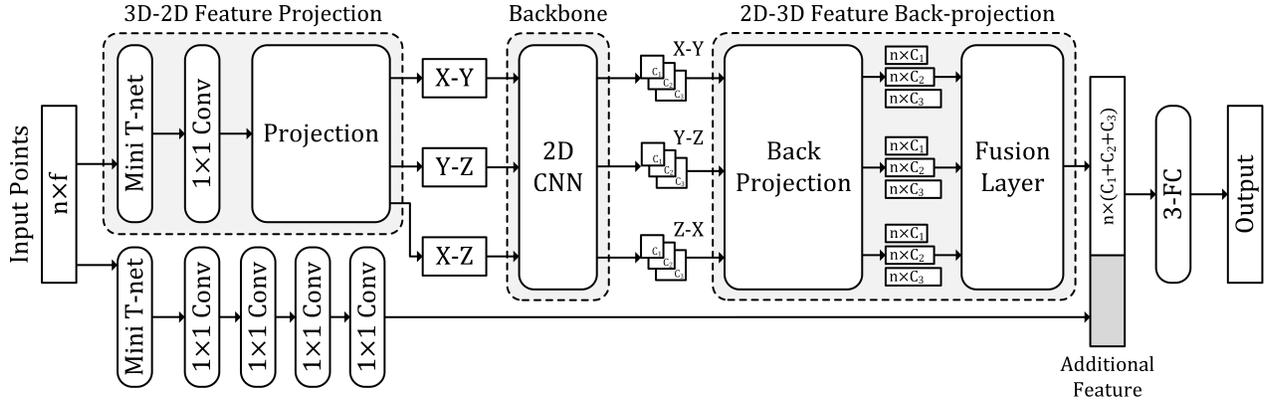


Fig. 2. Overview of our framework

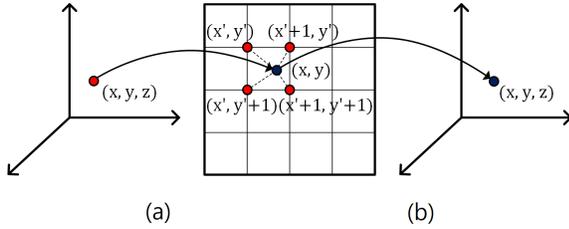


Fig. 3. The operation of the (a) 3D-to-2D feature projection module and the (b) 2D-to-3D feature back projection is illustrated

stage applies deconvolution to produce pixel-wise classification results from the global features. [24] reported a similar structure with [18] and also combined intermediate features with the same spatial resolution at two stages. In this paper, we only report experimental results when using FCN [16] as a baseline. However, it should be noted that other network architectures such as [18], [24] can be used for 2D feature extraction and can possibly improve the performance.

### III. METHOD

#### A. Overview

To demonstrate 3D feature extraction using 2D CNNs, we propose a framework that leverages 2D convolutional layers as a backbone. For input 3D point clouds, our architecture initially transforms them into 2D feature maps, then extracts hierarchical features using 2D convolutional layers, and finally projects the extracted features back onto 3D point locations. These three processes are separately explained in the following subsections. Our final architecture combines the proposed projection-based features with additional point-wise features to further improve performance, as explained in Section III-C. The overall architecture is presented in Figure 2.

#### B. Proposed Modules

1) *3D-to-2D Feature Projection*: The proposed architecture uses only point clouds as input, which are processed

using 2D convolutional layers. However, as point clouds lie in a continuous 3D space, they cannot be directly fed into 2D CNNs. Hence, the point clouds need to be converted into 2D feature maps that lie on a regular grid, so that they can be processed by 2D CNNs. Alignment and shallow 3D feature extraction of the input point clouds are performed using mini T-net followed by a single convolutional layer. Mini T-net is a translation network that was used in Pointnet [20]; it predicts a transformation matrix through a small network. We reduce the channel size of MLP layers to increase efficiency. Then, these point features are projected onto the XY-, YZ-, and ZX-planes.

Let  $p$  denote a single point in point clouds, with the coordinates  $(x, y, z)$  and the feature  $f$ . First, we normalize the point cloud so that it lies in the given range for a plane. If the given plane is the XY-plane and the plane length is  $l$ , then  $x$  and  $y$  are normalized to be within  $[-\frac{l}{2}, \frac{l}{2}]$ . Therefore, we put  $f$  on a grid on the XY-plane via bilinear interpolation. Although there are many interpolation methods that can be used to fill the grid, bilinear interpolation is adopted because of its low computation and memory requirement. If multiple features are assigned to the same cell of the grid, these features are added. This results in a 2D feature map on the  $l * l$  grid that is transformed from a 3D point cloud. This transformation can be represented as the following equation:

$$I_{XY}(x', y') = \sum_x \sum_y G(x', y', x, y) \cdot f(x, y, z) \quad (1)$$

$$G(A, B, a, b) = g(A, a) \cdot g(B, b) \quad (2)$$

$$g(N, n) = \max(0, 1 - |N - n|) \quad (3)$$

$I_{XY}(x', y')$  and  $f(x, y, z)$  denote the 2D features on position  $(x', y')$  of XY-plane and the 3D features of the point with coordinate  $(x, y, z)$  respectively.  $G(\cdot, \cdot)$  is the 2D bilinear interpolation kernel that is composed of two one-dimensional kernels, as shown in Equation 2. We present Figure 3 to help understand.

In the proposed method, three orthogonal planes are used for the projection rather than a single plane. The general projection and back-projection operations do not consider

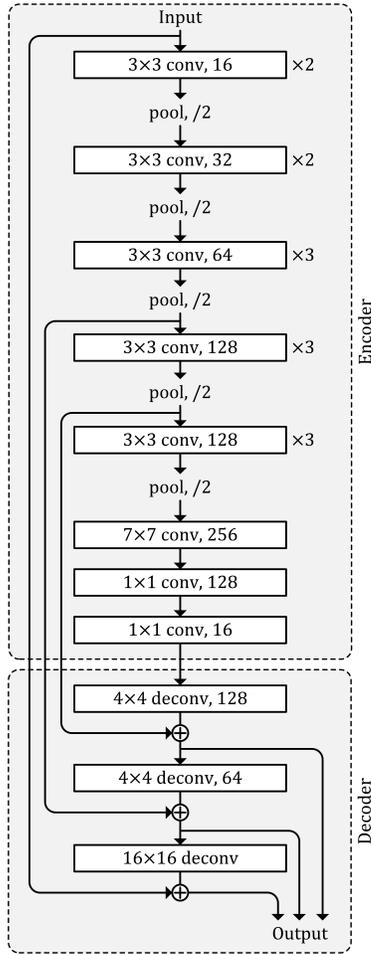


Fig. 4. The architecture of the modified FCN-VGG16

the information regarding the direction of projection; hence, using one or two planes does not guarantee that each point will possess distinctive features. Consider that there are two points  $p$  and  $p'$  with coordinates  $(x, y, z)$  and  $(x, y, z')$ , respectively. When the XY-plane is used, projection and back-projection are performed by only using  $x$  and  $y$  values. Therefore,  $p$  and  $p'$  will possess identical features. To overcome this issue, we employ three non-parallel planes for the projection; in particular the orthogonal planes are used for simplicity. This allows each projection to extract different sub-features that complement each other, thereby ensuring that the proposed module is able to generate distinctive features for each point by using 2D convolutions. We demonstrate the effectiveness of using three orthogonal planes through comparative experiment, as described in Section IV-C.

2) *2D Convolutional Backbone*: The 2D features that are transformed by above-mentioned projection are inputted to the 2D convolutional backbone network to extract high-level features. We use the backbone network based on the FCN-VGG16 [16], which has been widely used for 2D segmentation tasks. However, we implement two modifications to improve efficiency and performance. First, we reduce the size

of the channels in multiple layers by a factor of four to reduce the required memory and speed up the processing. Second, this network outputs feature maps from the intermediate layers as well as the final layer. This facilitates the creation of multi-scale features that possess rich information. We extract two intermediate feature maps that contain 128 and 64 channels and one final feature map that contains 16 channels from one input. As there are three inputs from the 3D-to-2D feature projection module, the total number of outputted feature maps is 9. This structure is presented in Figure 4.

3) *2D-to-3D Feature Back-projection*: To apply the 2D features extracted using the 2D convolutional backbone to the point clouds, we back-project the 2D features to the 3D points. This is the reverse operation of the method described in Section III-B.1. We normalize  $p$  to lie in the given image length range for a specific plane. Subsequently, we extract a new feature  $f'$  from the grid by using bilinear interpolation. Therefore, in the case of the XY-plane, the transformation can be represented as follows:

$$f_{XY}(x, y, z) = \sum_X \sum_Y G(x, y, x', y') \cdot I_{XY}(x', y') \quad (4)$$

$f_{XY}(x, y, z)$  and  $I_{XY}(x', y')$  denote the 3D features of the point with coordinate  $(x, y, z)$  from XY-plane and the 2D features at position  $(x', y')$  of the XY-plane, respectively.  $G(\cdot, \cdot)$  is the 2D bilinear interpolation kernel, equivalent to that in Equation 2. The operation is shown in Figure 3.

After back-projection, we obtain 9 sub-features that were extracted from the 2D convolutional network for each point cloud. We combine these 9 sub-features by using the fusion layer to create a single integrated feature. Let each feature be  $f_P^d$ , wherein  $P$  and  $d$  are the given projection plane and the number of output dimension, respectively.

The fusion layer first adds the sub-feature of different planes with the same dimension to create the intermediate feature  $f^d$  for each dimension.

$$f^{128} = f_{XY}^{128} + f_{YZ}^{128} + f_{XZ}^{128} \quad (5)$$

$$f^{64} = f_{XY}^{64} + f_{YZ}^{64} + f_{XZ}^{64} \quad (6)$$

$$f^{16} = f_{XY}^{16} + f_{YZ}^{16} + f_{XZ}^{16} \quad (7)$$

Subsequently, the integrated feature  $f_{final}$  is created via channel-wise concatenate of the intermediate features

$$f_{final} = (f^{128}, f^{64}, f^{16}) \quad (8)$$

### C. Architecture and loss

Our proposed architecture is simple and easy to implement. The overall architecture is composed of the three modules described in the previous subsection, followed by fully connected layers. For the final architecture, we add point-wise features to improve training. The additional features are simply extracted from the input point clouds using mini T-Net and fully connected layers. As every module

performs differentiable operations, our framework is end-to-end trainable. For the loss function, simple point-wise cross-entropy is used for training. The point-wise cross-entropy function is the same as pixel-wise cross-entropy used in image segmentation:

$$CE(q, p) = - \sum_n \sum_i q_{n,i} \log(p_{n,i}) \quad (9)$$

where  $p$  and  $q$  denote the predictions and ground-truth labels, respectively;  $n$  is the index of points, and  $i$  refers to the class index.

#### IV. EXPERIMENTS

To demonstrate the effectiveness of 2D convolutions for 3D point cloud segmentation tasks, we perform extensive experiments. First, we evaluate our proposed architecture on various datasets and compare it to other state-of-the-art methods. Thereafter, ablation studies are performed to further investigate the proposed modules.

##### A. Datasets and Metrics

This section summarizes the datasets and metrics used in experiments. We use S3DIS [2] and ShapeNet-Part [5] for segmentation task.

- **S3DIS.** This dataset contains 3D scans of 6 indoor areas including 271 rooms. Each point is acquired from Matterport scanners and annotated with one of the 13 classes.
- **ShapeNet-Part.** This dataset contains 16,880 models-14,006 for training and 2,874 for testing. Each point is annotated with 2 to 6 parts. This dataset contains 16 classes with 50 different parts in total.

The mean intersection over union (mIoU) was used as the evaluation metric of the segmentation tasks. The IoU is computed as  $\frac{TP}{TP+T-P}$ , where  $TP$ ,  $T$ , and  $P$  are the number of true positive points, number of groundtruth points, and the number of predicted positive points. The mIoU is calculated by averaging the IoU of each class.

##### B. Implementation

We implement the proposed architecture by using the Tensorflow framework [1]. The Adam optimizer is used and the size of each projection plane is set to (224, 224). To ensure that the sole performance of the 2D convolution operation could be analyzed, post processing techniques were not used. We emphasize that all the operations in the proposed architecture is implemented by using the built-in functions of Tensorflow and can be trained end-to-end from scratch.

- **S3DIS.** We sample 4,096 points from the dataset and train the network for 50 epochs. The initial learning rate is 0.001, and the decay rate 0.5. The learning rate decay step is 300,000.
- **ShapeNet-Part.** 2,048 points are sampled from the dataset. We train the network for 200 epochs, and set the initial learning rate as 0.001. The learning rate is divided by 2 after 200,000 steps.

TABLE I

THE RESULTS OF THE ABLATION STUDIES ON S3DIS AREA5 DATASET

Ablation	mIoU	# Planes	T-Net	Multi-Scale	Additional
# Planes	42.4	1			
	43.1	2			
	48.2	3			
T-Net	52.4	3	✓		
Multi-scale	53.2	3	✓	✓	
Additional	<b>56.8</b>	3	✓	✓	✓

##### C. Analysis

We investigate the performance of different architectures to analyze our architecture through the ablation studies. As described in Section III-B.1, We compare the performance of the projection module by subtracting orthogonal planes. We also study the effect of the transformation network that exists in front of the projection module, multi-scale feature maps from the 2D backbone network, and additional features of the final architecture. All the experiments are performed as segmentation tasks using the S3DIS dataset. The S3DIS Area 5 dataset is used for testing, and the remaining areas datasets are used for training.

**Projection of three orthogonal planes.** We vary the number of orthogonal planes on the projection module. Our experiments, which are listed in Table I (# Planes), prove that additional orthogonal planes facilitate the extraction of additional various features that complement each other. This is analogous to the approach used in multi-view-based studies [25], [28]. When the number of orthogonal planes in the projection are reduced, there is significant degradation in the performance. For two orthogonal planes and one orthogonal plane, the performance decreases by 5.1% and 5.8%, respectively.

**Transformation Network.** Table I (T-Net) lists the results of removing the transformation network that is in front of the projection module. We observe that the performance reduces by 4.2%, which indicates that the transformation network assists in training by aligning the input coordinates and additional data.

**Multi-scale feature maps.** We eliminated the multi-scale feature maps from the 2D backbone network. The result in Table I (Multi-Scale) shows that the performance decreases by 0.8%, which indicates that the multi-scale feature maps from the backbone network facilitate the creation of richer information by considering multi-scale features. This is inspired by the feature pyramid concept [15], which is widely used for 2D segmentation applications.

**Additional features.** The results in Table I (Additional) shows that the final architecture, which is added additional features from the baseline architecture, accounts for an improvement of 3.6%. This indicates that the additional features are well fused with the extracted features to assist in 3D feature extraction. It is noteworthy that even simple additional features such as those used in this study can improve performance.

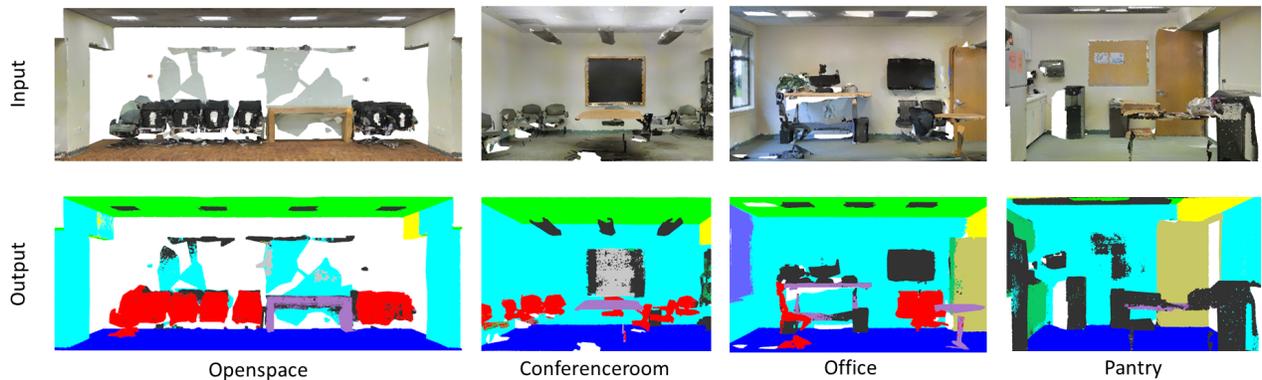


Fig. 5. Indoor scene semantic segmentation on the S3DIS dataset

TABLE II  
OBJECT PART SEGMENTATION SCORES ON THE SHAPENET-PART DATASET (MIDDLE), AND INDOOR SCENE SEMANTIC SEGMENTATION SCORES ON THE S3DIS DATASET (RIGHT)

Method	ShapeNet mIoU	S3DIS mIoU
SyncCNN [33]	82.0	-
SpiderCNN [32]	81.7	-
SplatNet [27]	<b>83.7</b>	-
SO-Net [13]	81.0	-
SPGN [29]	82.8	50.4
KCNet [26]	82.2	-
KdNet [8]	77.4	-
3DmFV-Net [4]	81.0	-
DGCNN [30]	82.3	56.1
RSNet [7]	81.4	56.5
PointNet [20]	80.4	47.6
PointNet++ [22]	81.9	-
SPG [10]	-	62.1
PointConv [31]	82.8	-
ResGCN-28 [12]	-	60.0
ours	82.7	<b>62.8</b>

#### D. Comparison

To verify the utility of 3D feature representations by using 2D convolution operations, we compare our final architecture to several state-of-the-art methods. We perform object part segmentation and semantic scene segmentation using ShapeNet-Part dataset and S3DIS dataset respectively. Unlike our previous ablation studies, we follow the 6-fold cross validation strategy for the S3DIS dataset. Table II clearly indicate that our architecture is comparable to many existing methods, which means that 3D feature representation via the 2D convolution operation is effective, in terms of performance accuracy.

The results in Table II (middle) show that the proposed architecture has a comparable performance with state-of-the-art methods, for object part segmentation task. Qualitative comparisons between our results and the groundtruth are depicted in Figure 6. We observe that the proposed method performs part segmentation robustly on many objects in the ShapeNet-Part dataset.

Table II (right) shows the results of the indoor scene semantic segmentation task. Our architecture outperforms

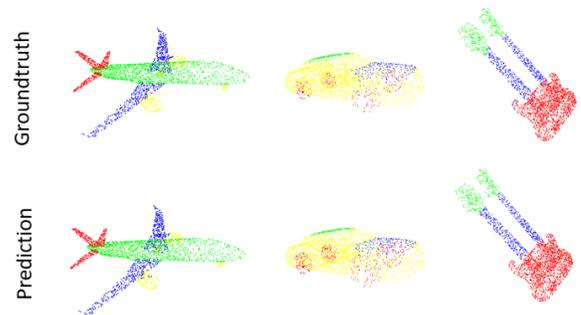


Fig. 6. Object part segmentation on the ShapeNet-Part dataset

most state-of-the-art methods, and in particular, we can see that it outperforms [10] by 0.7%. We also visualize our results and the groundtruth in Figure 5. It can be seen that the proposed architecture performs fairly well in general cases.

We demonstrate the comparable results with state-of-the-art methods and believe that adjusting the hyper-parameters for training can result in further improvement.

#### V. CONCLUSION

In this paper, we propose PBP-Net, which is a simple yet effective neural network architecture leveraging 2D CNNs for 3D point cloud segmentation. In order to handle 3D point clouds, PBP-Net first projects 3D point cloud and their local features on multiple 2D planes. The features extracted by the 2D convolution backbone are re-allocated to their original point locations. We prove that a 2D convolution-based method can show comparable performance to that of other state-of-the-art methods.

PBP-Net provides a new perspective on the projection-based method as well as a baseline for using 2D CNNs for 3D tasks. Furthermore, the proposed method can be improved with the development of 2D segmentation, since we do not make any modification to the 2D CNNs when using them as backbone. We expect that this research will lead to the integration of 2D vision studies to 3D technologies, which have been remaining independent thus far.

## ACKNOWLEDGEMENTS

This research was supported by Naver Labs Corporation.

## REFERENCES

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI16), pages 265–283, 2016.
- [2] Iro Armeni, Ozan Sener, Amir R. Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1534–1543, 2016.
- [3] Jens Behley, Martin Garbade, Andres Milioto, Jan Quen-zel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In Proceedings of the IEEE International Conference on Computer Vision, pages 9297–9307, 2019.
- [4] Yizhak Ben-Shabat, Michael Lindenbaum, and Anath Fischer. 3dmfv: Three-dimensional point cloud classification in real-time using convolutional neural networks. IEEE Robotics and Automation Letters, 3(4):3145–3152, 2018.
- [5] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012, 2015.
- [6] Binh-Son Hua, Minh-Khoi Tran, and Sai-Kit Yeung. Pointwise convolutional neural networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 984–993, 2018.
- [7] Qiangui Huang, Weyue Wang, and Ulrich Neumann. Recurrent slice networks for 3d segmentation of point clouds. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2626–2635, 2018.
- [8] Roman Klokov and Victor Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In Proceedings of the IEEE International Conference on Computer Vision, pages 863–872, 2017.
- [9] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.
- [10] Loïc Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4558–4567, 2018.
- [11] Alex H. Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 12697–12705, 2019.
- [12] Guohao Li, Matthias Müller, Ali Thabet, and Bernard Ghanem. Deepgens: Can gcnns go as deep as cnns? In Proceedings of the IEEE International Conference on Computer Vision, pages 9267–9276, 2019.
- [13] Jiaxin Li, Ben M. Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 9397–9406, 2018.
- [14] Yangyan Li, Sören Pirk, Hao Su, Charles R. Qi, and Leonidas J. Guibas. Fpnn: Field probing neural networks for 3d data. In Advances in Neural Information Processing Systems, pages 307–315, 2016.
- [15] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2117–2125, 2017.
- [16] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 3431–3440, 2015.
- [17] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 922–928. IEEE, 2015.
- [18] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In Proceedings of the IEEE international conference on computer vision, pages 1520–1528, 2015.
- [19] Jeremie Papon, Alexey Abramov, Markus Schoeler, and Florentin Worgatter. Voxel cloud connectivity segmentation - supervoxels for point clouds. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2027–2034, 2013.
- [20] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 652–660, 2017.
- [21] Charles R. Qi, Hao Su, Matthias Nießner, Angela Dai, Mengyuan Yan, and Leonidas J Guibas. Volumetric and multi-view cnns for object classification on 3d data. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 5648–5656, 2016.
- [22] Charles R. Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In Advances in neural information processing systems, pages 5099–5108, 2017.
- [23] Dario Reithage, Johanna Wald, Jurgen Sturm, Nassir Navab, and Federico Tombari. Fully-convolutional point networks for large-scale point clouds. In Proceedings of the European Conference on Computer Vision (ECCV), pages 596–611, 2018.
- [24] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical image computing and computer-assisted intervention, pages 234–241. Springer, 2015.
- [25] Riccardo Roveri, Lukas Rahmann, Cengiz Oztireli, and Markus Gross. A network architecture for point cloud classification via automatic depth images generation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4176–4184, 2018.
- [26] Yiru Shen, Chen Feng, Yaoqing Yang, and Dong Tian. Mining point cloud local structures by kernel correlation and graph pooling. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4548–4557, 2018.
- [27] Hang Su, Varun Jampani, Deqing Sun, Subhansu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. Splatnet: Sparse lattice networks for point cloud processing. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2530–2539, 2018.
- [28] Hang Su, Subhansu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In Proceedings of the IEEE international conference on computer vision, pages 945–953, 2015.
- [29] Weyue Wang, Ronald Yu, Qiangui Huang, and Ulrich Neumann. Sgpn: Similarity group proposal network for 3d point cloud instance segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2569–2578, 2018.
- [30] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. ACM Transactions on Graphics (TOG), 38(5):146, 2019.
- [31] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 9621–9630, 2019.
- [32] Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. Spiderncnn: Deep learning on point sets with parameterized convolutional filters. In Proceedings of the European Conference on Computer Vision (ECCV), pages 87–102, 2018.
- [33] Li Yi, Hao Su, Xingwen Guo, and Leonidas J. Guibas. SyncSpecCnn: Synchronized spectral cnn for 3d shape segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2282–2290, 2017.
- [34] Chris Zhang, Wenjie Luo, and Raquel Urtasun. Efficient convolutions for real-time semantic segmentation of 3d point clouds. 2018 International Conference on 3D Vision (3DV). IEEE, 2018.