MARNet: Multi-Abstraction Refinement Network for 3D Point Cloud Analysis

Rahul Chakwate   Arulkumar Subramaniam   Anurag Mittal
Indian Institute of Technology Madras
ae16b005@smail.iitm.ac.in

Abstract

Representation learning from 3D point clouds is challenging due to their inherent nature of permutation invariance and irregular distribution in space. Existing deep learning methods follow a hierarchical feature extraction paradigm in which high-level abstract features are derived from low-level features. However, they fail to exploit different granularity of information due to the limited interaction between these features. To this end, we propose Multi-Abstraction Refinement Network (MARNet) that ensures an effective exchange of information between multi-level features to gain local and global contextual cues while effectively preserving them till the final layer. We empirically show the effectiveness of MARNet in terms of state-of-the-art results on two challenging tasks: Shape classification and Coarse-to-fine grained semantic segmentation. MARNet significantly improves the classification performance by 2% over the baseline and outperforms the state-of-the-art methods on semantic segmentation task.

1. Introduction

The recent evolution of 3D sensors such as LiDAR and Kinect has boosted the ability to perceive the environment in terms of 3D point clouds. 3D Point cloud analysis has increasingly become ubiquitous in robotic perception [36], augmented / mixed reality [16] and autonomous driving [64, 9, 48]. Traditional analysis techniques with hand-crafted features are unable to cater to the needs of these modern applications that demand diverse semantic understanding. With a proven track record in terms of state-of-the-art solutions in multiple domains such as image analysis [18, 46, 10, 14, 49] and natural language processing [51, 1, 34], deep learning pledges a promising alternative to conventional methods. Owing to the remarkable improvements of deep neural network architectures in 2D image analysis, researchers attempt to port and adapt the 2D network architectures to 3D point clouds [38, 40, 25, 27].

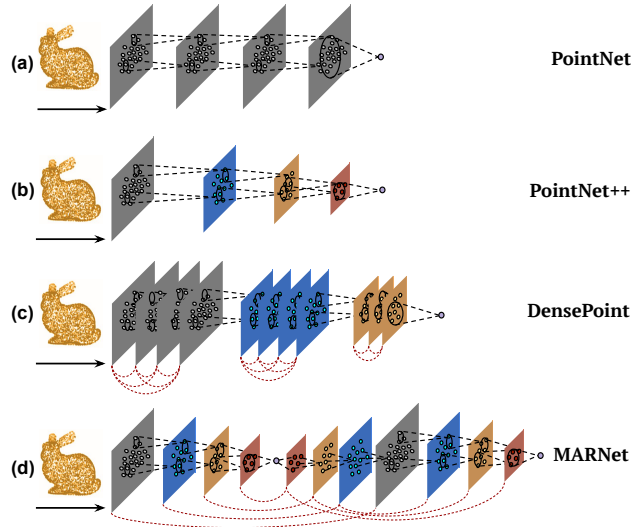Code available at: https://github.com/ruc98/MARNet

Figure 1: State-of-the-art models for feature extraction from 3D point clouds. (a) PointNet [38] uses a series of point-wise MLPs followed by max/average pooling, (b) Point-Net++ [40] introduces the notion of hierarchical feature extraction to capture local patterns, (c) DensePoint [27] uses densely connected blocks to aggregate contextual information, (d) **MARNet (ours)** proposes a unified architecture to perform hierarchical and multi-level feature aggregation to capture local and global contextual cues.

A natural perspective is to convert irregular point clouds to regular grid voxels [32, 57, 20, 4] and apply 3D CNNs to extract features from them. However, voxelization may obscure the finer details due to lower grid resolution. On the other hand, the cost of making the voxel grids finer is exponentially high. Another strategy is to project the 3D information on multi-view images [47, 63, 62, 39]. However, it suffers from the loss of crucial 3D geometric information and thus lack rich contextual features.

A pioneering approach called PointNet [38] proposes to directly operate on irregular point clouds. It transforms the 3D points using a series of point-wise feature transformations and finally, outputs an aggregated feature vector using a permutation invariant symmetric function. This successful method, however, has a downside that local patterns are not taken into account. To mitigate this issue, PointNet++ [40]

1

proposes to group the neighboring points in the euclidean space and apply PointNet [38] locally on each of the group, thus inducing the notion of hierarchical feature extraction. This architectural design resembles the hierarchical feature extraction of CNNs used for 2D image analysis [18].

In their solution to incorporate local patterns, Point-Net++ [40] proposes two layers, namely: 1) Multi-Scale Grouping (MSG) layer and 2) Multi-Resolution Grouping (MRG) layer. The MSG layer aggregates point-wise features at different scales (i.e., group the points with multiple radii). Whereas, MRG layer aggregates the point features at different resolutions (i.e., from multiple abstraction layers). Strong empirical performance of [40] suggests that both these layers aid in capturing local patterns. However, these layers fail to apprehend dense contextual insights, as there is a lack of feature interaction between the global features in deeper layers and local features in earlier layers. Inspired by architectural improvement in the 2D image analysis domain [14], DensePoint [27] tries to diminish this issue by using densely-connected blocks to encourage feature reuse and enhance feature propagation in 3D point clouds. On the downside, it fails to support multi-level feature interaction and lacks in its ability to preserve the features from all stages, as the earlier features are modified during the course of forward propagation.

As a result of the above observations, to achieve optimal performance in the point cloud analysis tasks, we notice that the following requirements are crucial in the network: 1) Multi-scale and multi-resolution aggregation of the features to capture local as well as global patterns, 2) Efficient feature communication between the shallow and deeper features to gain dense-contextual information, 3) Preservation of features at all levels of abstraction for effective feature learning and back-propagation.

To this end, we propose a novel deep network architecture for 3D point clouds analysis: **M**ulti-**A**bstraction **R**efinement **N**etwork (MARNet), by carefully designing the network layers to satisfy the above requirements. Precisely, MARNet (Fig. 1) consists of three stages in its network design, namely: 1) Backbone stage, 2) Feature Cross-Referencing (FCR) stage, and 3) Feature Re-Encoding (FRE) stage. First, the Backbone stage employs a hierarchical multi-scale feature extractor such as PointNet++ [40], thus incorporating multi-scale and multi-resolution feature learning. Next, to encourage effective feature propagation between shallow and deep layers, the FCR stage fosters an efficient interaction of the multiple abstraction features from the Backbone stage and allows them to refine each other using a specially devised parameter-less reduction function. Further, FCR and FRE stages are designed to preserve features from all the Backbone levels by incorporating residual connections until the final output layer, thus encouraging unimpeded gradient flow. Our three-stage de-

sign of MARNet is inspired by one of the recent deep learning architectural improvements in the 2D image analysis domain called FishNet [49] that makes use of multi-level feature interaction to achieve state-of-the-art performance in 2D image classification, detection, and segmentation.

The key contributions of our paper are as follows:

- We propose a novel unified framework that utilizes the complementary nature of multi-level abstract features and encourages them to interact and refine each other.

- We carefully design network layers to preserve point-features of different granularity such that unmodified gradients are passed to earlier layers to overcome vanishing/exploding gradient problem

- Through extensive experiments, we verify the effectiveness of MARNet by attaining state-of-the-art results on challenging benchmarks for the tasks: 3D shape classification, coarse-, middle- and fine-grained semantic segmentation.

2. Related Work

We delineate the existing works on 3D point cloud analysis into the following categories:

Multi-view projection-based and volumetric methods: The multi-view projection-based methods [47, 63, 62, 39, 7, 52, 29] project a 3D point cloud onto multiple 2D views and apply 2D CNN to extract view based features. These multi-view features are aggregated to get a global representation of the shape. Naively projecting 3D objects onto 2D space leads to loss of valuable geometric information. Other methods [32, 57, 20, 4, 43, 8] project 3D point cloud onto regular 3D grid voxel. However, voxelization leads to quantization loss caused by the low grid resolution of the voxels. The computation increases exponentially with a linear increase in grid resolution. Kd trees [17] and octree [22, 41, 54] based methods alleviate these limitations though they are still dependent on subdivision of volume. On the contrary, our model learns directly from irregular point clouds.

Point-wise MLP based Methods: These types of networks operate directly on each point. PointNet [38] applies a shared-MLP on each point independently and aggregates the point features using max-pooling to obtain a global representation. However, it fails to capture the local patterns as the features are learned independently for every point. PointNet++ [40] overcomes this limitation by hierarchically grouping and downsampling the point clouds. Many subsequent networks [61, 5, 13, 19, 27] including ours are based on these networks.

Convolutional Kernel-Based Methods: Kernel-based networks also involve point-wise MLPs. However, they have specialized 3D convolutional kernels that operate either locally or globally on the point clouds. [56, 11, 60, 31]
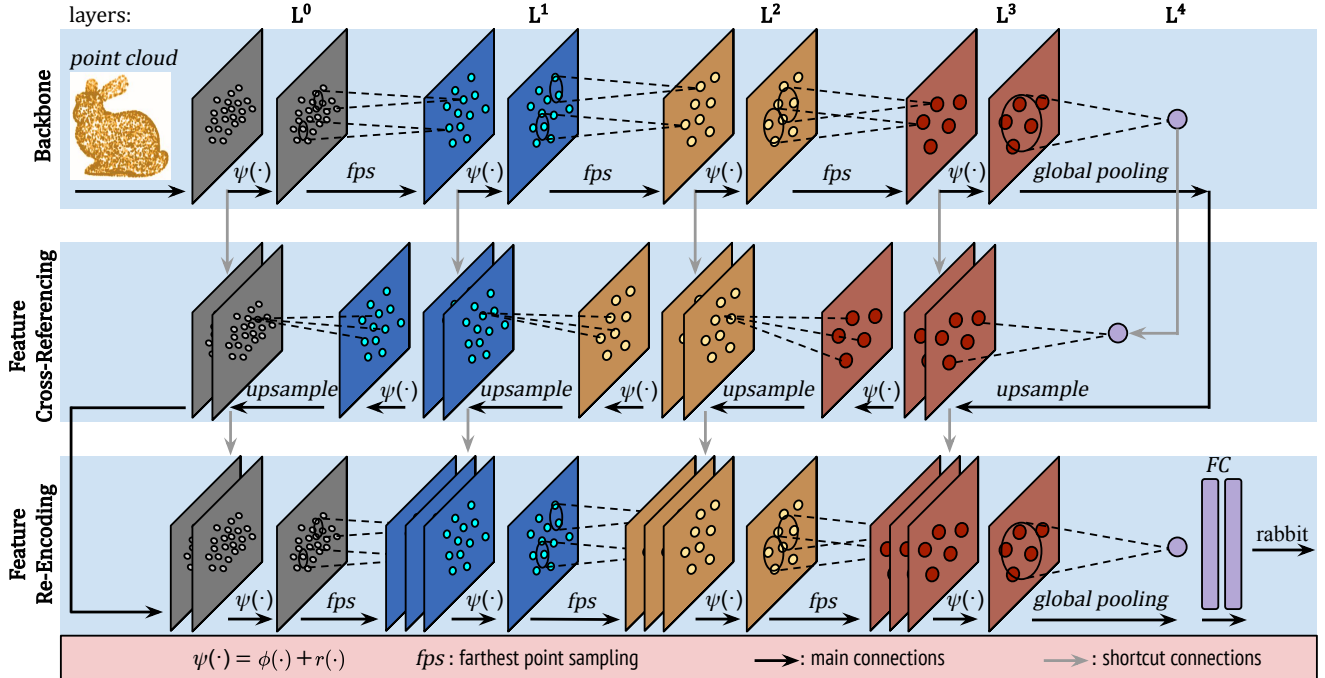
Figure 2: The overall network architecture containing three stages of feature extraction/refinement: 1) Backbone stage (top), 2) Feature Cross-Referencing stage (middle) and 3) Feature Re-Encoding stage (bottom). $\psi(\cdot)$ = transformation function $\phi(\cdot)$+ residual function $r(\cdot)$ as mentioned in Sec. 3. In both Backbone & FRE stage, $r(\cdot)$ is implemented as identity function, whereas in FCR stage, $r(\cdot)$ is implemented as a reduction function (Sec. 3.3).

use existing methods such as Monte Carlo estimation, Taylor expansion, k-nearest neighbors to model the 3D convolution operation. [21, 31, 37] address the rotation equivariance of point clouds by transforming the problem into polar coordinates. Several methods [19, 25, 30] cornerstone the point clouds' inherent geometric aspects while [27, 59, 28] focus on context aggregation. In contrast, our method focuses on preserving and refining the features obtained from the points and making the best use of the available features.

Graph Based Methods: These networks model point clouds as a graph with each point representing the vertex and the connection between the points as edges. [45, 44, 55, 65] model the distance between the points as the weights of the edges to capture the local geometry of the points. [26, 3] aim at easing the task of point agglomeration into simple steps. [50, 24, 6, 53] exploit the spectral domains for graph creation.

3. Methodology

Our network architecture consists of three stages of feature extraction and refinement, namely: 1) Backbone stage, 2) Feature Cross-Referencing stage (FCR) stage, and 3) Feature Re-Encoding (FRE) stage. The proposed network architecture is shown in Fig. 2.

The stages of our network are explained in Sec. 3.2, Sec. 3.3 and Sec. 3.4 respectively. The unique features of MAR-Net, such as multi-level feature aggregation and unimpeded

backpropagation of gradients, are discussed in Sec. 3.5.

3.1. Mathematical Notation

The different "levels" of abstraction are labeled as L^i where $i \in [0, 1, \dots]$ denotes the level index. The collection of 3D points at level L^i is given as $P^i = [p_1, p_2, \dots, p_{N^i}] \in \mathbb{R}^{N^i \times 3}$ where N^i is the number of 3D points at level L^i and dimensions of 3D points are $[x, y, z]$. The collection of point-wise features at level L^i is given by $F^i = [f(p_1), f(p_2), \dots, f(p_{N^i})] \in \mathbb{R}^{N^i \times D^i}$, where $f(p_i) \in \mathbb{R}^{D^i}$ denotes the point-wise feature vector corresponding to the 3D point p_i, D^i = feature dimension at level L^i.

Any entity $\mathbb{X}$ (such as Level L^i, 3D points p^i, point features f^i) corresponding to different stages (Backbone, FCE, FRE) are denoted by subscripts: $\mathbb{X}_{bb}$ for Backbone, $\mathbb{X}_{fcr}$ for FCR stage and $\mathbb{X}_{fre}$ for FRE stage. For example, L_{bb}^i denotes i^{th} level of Backbone stage.

3.2. The Backbone Stage

The Backbone stage has l levels denoted as $\{L_{bb}^0, L_{bb}^1, \dots, L_{bb}^{l-1}\}$. At every level L_{bb}^i, corresponding 3D points P_{bb}^i (along with its features F_{bb}^i) are hierarchically down-sampled and passed to the next level L_{bb}^{i+1}, i.e., $N_{bb}^i > N_{bb}^{i+1}$. We use the existing network PointNet++ [40] as our backbone for its simpler design, multi-level nature of feature extraction and widespread use. However, other hierarchical feature extractors [27, 25, 55]

3

can also be used as backbone. Specifically, we use multi-Scale Grouping (MSG) version of PointNet++ [40] to capture local neighborhood information at different scales.

First, we briefly introduce our backbone (Point-Net++ [40]) and then mention the key modifications made by us for efficient processing. At every level L_{bb}^i of our backbone, first, N_{bb}^{i+1} center points are sampled using farthest point sampling (FPS) [40] technique. Points within a sphere of radius d, centered on these points are selected and a shared local PointNet [38] aggregates their features. Multiple radii are used to capture local shape variations and contextual cues. The mathematical formulation of grouping and feature aggregation can be written as:

$$f_{bb}^{i+1}(p) = \sigma(\{\phi(f_{bb}^i(q)); \forall\, q \in ||q - p||_2 \leq d\}) \quad (1)$$

where $p \in P_{bb}^{i+1}$ is the center point, $q \in P_{bb}^i$ is within the distance d from p, $\phi(\cdot)$ is a shared point-wise feature transformation layer (shared-MLPs) and $\sigma(\cdot)$ is a symmetric aggregation function (a max pooling layer).

Next, we perform two modifications in the Backbone for reducing parameters as well as for efficient gradient propagation: 1) To reduce the parameters, we replace feature transformation layers (MLPs) with grouped convolutions [18] inspired by [27], where the input channels are divided into N_g groups and convolved separately. 2) To facilitate direct gradient propagation in the backbone network, we add residual connections [10] across the grouped convolution layers. Eqn. 1 is modified to the utilize residual connections as follows:

$$f_{bb}^{i+1}(p) = \sigma(\{\phi(f_{bb}^i(q)) + r(f_{bb}^i(q)); \\ \forall\, q \in ||q - p||_2 \leq d\}) \quad (2)$$

Henceforth, $\phi(\cdot)$ denotes a point-wise feature transformation function with grouped convolutions and $r(\cdot)$ denotes a residual function unless otherwise specified. In the Backbone stage, $r(\cdot)$ is implemented as an identity function.

In the final backbone level (L_{bb}^{l-1}), the features are aggregated using a global max pooling function and passed to the FCR stage.

3.3. Feature Cross-Referencing (FCR) Stage

The motivation of this stage is to preserve and refine the features from backbone by letting the low-level and high-level abstract features interact with each other. Such interactions between multi-level abstract features are proven crucial in segmentation [42] as well as recognition [49] tasks. To achieve this, we carefully design learning blocks to merge multi-level abstract features.

Let $\{F_{bb}^0, F_{bb}^1, \ldots, F_{bb}^{l-1}\}$ be the features from $\{L_{bb}^0, L_{bb}^1, \ldots, L_{bb}^{l-1}\}$ levels of Backbone stage. As Backbone stage has l levels of hierarchy to extract multi-level abstract features, FCR stage also has l levels to merge


Figure 3: The functionality of level i in FCR stage. First, features from the same level of Backbone (f_{bb}^i) and FCR stage (f_{fcr}^i) are concatenated. Then, concatenated features are passed through a point-wise transformation ($\phi(\cdot)$) and reduction ($r(\cdot)$) layers. The output features from these two layers are added and further upsampled using 3NN interpolation [40]) technique to pass to next level.

and refine features from the Backbone. This design is similar to a decoder network in [42]. FCR levels are numbered from $l - 1$ to 0 such that $N_{fcr}^i = N_{bb}^i$, $N_{fcr}^i < N_{fcr}^{i-1}$, i.e., the number of points of Backbone and FCR stages at a particular abstraction level are kept equal. The final level backbone features F_{bb}^{l-1} are directly considered as F_{fcr}^{l-1}. Next, the functionality of every other FCR level can be defined in three steps (Fig. 3) as follows:

$$F_{cat}^i = concat(F_{fcr}^i, F_{bb}^i) \quad (3)$$
$$F_{ref}^i = \phi(F_{cat}^i) + r(F_{cat}^i) \quad (4)$$
$$F_{fcr}^{i-1} = upsample(F_{ref}^i) \quad (5)$$

In the first step (Eqn. 3), the point-wise feature vectors at the level i from Backbone and FCR stages are concatenated together (F_{cat}^i). Next, in the second step (Eqn. 4), to get refined features (F_{ref}^i) from concatenated features, F_{cat}^i is passed through a point-wise transformation function $\phi(\cdot)$ and a residual function $r(\cdot)$. Here, $r(\cdot)$ is implemented as a reduction function as explained below.

Reduction function $r(\cdot)$: It takes a feature vector f of D dimension as input and aggregates the features by summing up k adjacent feature dimensions of the feature vector. Thus the resulting feature dimension is $\lfloor \frac{D}{k} \rfloor$ (D is chosen to be divisible by k). This reduction function serves two purposes: 1) to reduce the feature dimension without involving new parameters, and 2) to backpropagate unmodified gradients to earlier layers. Mathematically, the reduction function $r(\cdot)$ can be written as:

$$r(f) = \left[\sum_{j=1}^{k} f[j], \sum_{j=k+1}^{2k} f[j], ..., \sum_{j=D-k}^{D} f[j] \right] \quad (6)$$

Here, $f[j]$ denotes j^{th} dimension of feature vector f.

The output features from two transformations ($\phi(\cdot), r(\cdot)$) are added channel-wise to form refined features F_{ref}^i. In the third step (Eqn. 5), the refined features are up-sampled to

4

Figure 4: The functionality of level i in FRE stage. The features from all three stages (Backbone, FCR, FRE) are concatenated and passed through a residual point-wise transformation layer ($\phi(\cdot) + r(\cdot)$). $r(\cdot)$ is an identity function. The output features are propagated to the next level by 'sampling, grouping, and feature averaging'.

the high resolution points at next level F_{fcr}^{i-1} by means of three nearest neighbor (3NN) interpolation method [40].

3.4. Feature Re-Encoding (FRE) Stage

In the pipeline so far, the multi-level features from the Backbone stage have been cross-referenced and merged by FCR stage to capture local and global contextual cues. Similar to the Backbone and FCR stages, FRE stage has l levels. In each of them, the multi-level features from both Backbone and FCR stages are combined and re-encoded to summarize multi-abstract features. The functionality of each FRE level is written in three steps (Fig. 4) as follows:

$$F_{cat}^i = concat(F_{fre}^i, F_{fcr}^i, F_{bb}^i) \qquad (7)$$

$$F_{ref}^i = \phi(F_{cat}^i) + r(F_{cat}^i) \qquad (8)$$

$$F_{fre}^{i+1} = downsample(F_{ref}^i) \qquad (9)$$

In the first step (Eqn. 7), features of the same abstraction level from all stages (Backbone, FCR, FRE) are concatenated together to get concatenated features F_{cat}^i. In the second step (Eqn. 8), concatenated features F_{cat}^i are passed through the point-wise shared-MLP layers ($\phi(\cdot)$) with grouped convolutions along with the residual function ($r(\cdot)$) to obtain the refined features F_{ref}^i. Here, similar to the Backbone stage, $r(\cdot)$ is implemented as an identity function. Further, in the third step (Eqn. 9), to pass N_{fre}^{i+1} points to $(i+1)^{th}$ level (where $N_{fre}^{i+1} < N_{fre}^i$), we select N_{fre}^{i+1} center points using the farthest point sampling technique and the 3D points are grouped in a Euclidean metric space around the selected centers. Then, the features corresponding to these points in each group are aggregated and passed to the next level. This process of grouping and aggregation is similar to the Backbone stage.

The final level features of FRE stage F_{fre}^{l-1} are used for predicting the target tasks (shape classification in Sec. 4.1 and semantic segmentation in Sec. 4.2). The specification of exact architectures for target tasks is given in the supplementary material.

3.5. Unique characteristics of MARNet

The two key characteristics that make our network unique are:

Multi-Abstraction Feature Aggregation: Our network aggregates the features of three granularity (from Backbone, FCR, FRE) in Re-encoding (FRE) stage. Hence, the final classification layer receives refined features from every level of abstraction. Further, this rich and diverse information can be used by the downstream classification/segmentation layers to precisely classify/segment the objects/parts.

Unimpeded Gradients Propagation: Exploding/Vanishing gradient [12] is a well-known problem in deep networks that makes the network training unstable. In our network, the corresponding abstraction levels in all three stages are connected to the final (classification/segmentation) layer through direct connections. For instance, Backbone features from level L_{bb}^i is merged to FCR stage's level L_{fcr}^i through feature concatenation (Eqn. 3 & 4) and reduction function, which does not involve any parameters. Similarly, FCR stage's features from level L_{fcr}^i are passed via FRE stage through a concatenation and identity residual layer (Eqn. 7 & 8) to the final layer. In this way, every layer of our network receives the gradients directly from the loss function in an unmodified way.

3.6. Implementation Details:

MARNet is implemented using Pytorch [35] framework. We optimize the models using Adam optimizer [15] with following hyperparameters: initial learning rate = 0.001, weight decay = 0.01, batch size = 32. The learning rate is decayed by multiplying with 0.7 after every 20 epochs. We train our model with an Nvidia GTX 1080Ti GPU. We elaborate on details about the models and training in the supplementary material.

4. Experiments

We validate the effectiveness of MARNet through rigorous experimentation on the tasks, namely: 1) Shape classification and 2) Coarse-to-fine grained semantic segmentation.

4.1. Shape Classification

Datasets and evaluation metrics: ModelNet40 and ModelNet10 [58] are well known benchmarks for 3D point cloud recognition. Both datasets contain objects in the form of 3D CAD models. ModelNet40 contains 9843 train objects and 2468 test objects from 40 different object categories. ModelNet10 has 3991 train samples and 908 test samples divided into 10 object categories. We evaluate the performance of our model with two evaluation metrics: 1) Overall Accuracy (*OA*) - the ratio number of correctly predicted objects to the total number of objects in the dataset,

Method	#points	OA (%)	mcA (%)
PointNet [38]	1k	89.2	86.2
SO-Net [23]	1k	89.4	-
SCN [59]	1k	90.0	87.6
Kd-Net(depth=10) [17]	1k	90.6	86.3
PointNet++ [40]	1k	90.7	-
Spec-GCN [53]	1k	91.8	-
DGCNN [55]	1k	92.2	90.2
PointCNN [25]	1k	92.2	88.1
PCNN [31]	1k	92.3	-
DensePoint [27]	1k	93.2	-
GeoCNN [19]	1k	93.4	**91.1**
Ours	**1k**	**93.9**	**91.1**
SO-Net [23]	2k	90.9	87.3
Kd-Net(depth=15) [17]	32k	91.8	88.5
PointNet++ [40]	5k	91.9	-
Spec-GCN [53]	2k	92.1	-
SpiderCNN [60]	5k	92.4	-
SO-Net [23]	5k	93.4	90.8

Table 1: Classification Results on ModelNet40 dataset. Here, *OA:* Overall Accuracy, *mcA:* mean class Accuracy. (best results are in **bold**)

Method	#points	OA (%)	mcA (%)
ECC [45]	1k	90.8	90.0
Kd-Net(depth=10) [17]	1k	93.3	92.8
KCNet [44]	1k	94.4	-
PCNN [31]	1k	94.9	-
DensePoint [27]	1k	**96.6**	-
Ours	**1k**	96.1	**95.9**
Kd-Net(depth=15) [17]	32k	94.0	93.5
SO-Net [23]	2k	94.1	93.9
SO-Net [23]	5k	95.7	95.5

Table 2: Classification results on ModelNet10 dataset. Here, *OA:* Overall Accuracy, *mcA:* mean class Accuracy. (best results are in **bold**)

2) mean class accuracy (*mcA*) - the average of the ratio of correctly predicted objects within a class to the total number of objects in the class. The *OA* metric measures the overall performance of the model, whereas *mcA* metric is more robust to the data imbalance within the classes.

Following [19, 40], during training, we uniformly sample 1024 points along with point normals as input to the model. To make a fair comparison with the state-of-the-art, we use the same data augmentation techniques used by [27, 17] to anisotropically scale the point clouds randomly in the range of [0.66, 1.5] and translate them in range [−0.2, 0.2]. During testing, we use voting evaluation as followed in [27, 38, 40] to average predictions from 10 runs.

Comparison with state-of-the-art methods: The performance comparison with the state-of-the-art [38, 23, 59, 17, 40, 53, 55, 25, 31, 27, 19] methods in the ModelNet40 dataset is shown in Tab. 1. For convenience, we categorize the methods based on the input number of points for

the model. First, we notice that MARNet significantly outperforms its backbone method PointNet++ [40], with an improvement of about +3% in overall accuracy. This improvement with a large margin validates our hypothesis that allowing multi-abstraction features to refine each other and preserving them from earlier layers is important for point cloud analysis. Further, our method also outperforms all the previous state-of-the-arts. Specifically, MARNet improves the performance over GeoCNN [19] and DensePoint [14] by +0.5% and +0.7% respectively on the overall accuracy. In terms of mean class accuracy (*mcA*), we perform on-par to GeoCNN [19] and outperform all other methods in ModelNet40 dataset. Experiments on ModelNet10 dataset (Tab. 2) shows that MARNet performs on-par with Dense-Point [27] and outperform all other state-of-the-arts in both the metrics (*OA* and *mcA*).

4.2. Coarse-to-fine grained Semantic Segmentation

Dataset and evaluation metric: PartNet [33] is a semantic (part) segmentation dataset consisting of three levels of shapes namely: 1) coarse, 2) middle, and 3) fine-grained. It is build on top of ShapeNet [2] and contains 26,671 3D models from 24 different object categories with 573,585 annotated part instances.

We evaluate our model on all three levels of segmentation. Following [33], we use part-category mIOU (in %) as the evaluation metric and compare our results with the published state-of-the-art methods: PointNet [38], PointNet++ [40], SpiderCNN [60], PointCNN [25]. In our experiments, we train individual networks for segmentation of different object categories and levels (coarse, middle and fine-grained), as followed in [33]. Precisely, we train 50 separate networks: 24 for coarse-level, 9 for middle-level and 17 for fine-grained level.

Comparison with state-of-the-art methods: Tab. 3 compares the performance of MARNet with other state-of-the-art methods on three different tasks: coarse-, middle- and fine-grained semantic segmentation of objects. The results are averaged across the three tasks as well as across the categories for each of the tasks. MARNet performs significantly better on the coarse-grained task with +1.8% improvement over the previous best performing method, PointNet++ [40]. On average, across all three tasks, our model performs better than the previous state-of-the-arts with +2.5% and +0.8% over PointNet++ [40] and PointCNN [25], respectively. The qualitative results shown in Fig. 5 illustrates the effectiveness of MARNet in semantic segmentation task.

5. Ablation Study

We conduct several ablation studies to verify the architectural choices adopted for MARNet, its robustness to the variations in input data and its computation efficiency.

Method	Avg	Bag	Bed	Bott	Bowl	Chair	Clock	Dish	Disp	Door	Ear	Fauc	Hat	Key	Knife	Lamp	Lap	Micro	Mug	Ref	Scis	Stora	Table	Trash	Vase
P1	57.9	42.5	32	33.8	58	64.6	33.2	76	86.8	64.4	53.2	58.6	55.9	65.6	62.2	29.7	96.5	49.4	80	49.6	86.4	51.9	50.5	55.2	54.7
P2	37.3	-	20.1	-	-	38.2	-	55.6	-	38.3	-	-	-	-	-	27	-	41.7	-	35.5	-	44.6	34.3	-	-
P3	35.6	-	13.4	29.5	-	27.8	28.4	48.9	76.5	30.4	33.4	47.6	-	-	32.9	18.9	-	37.2	-	33.5	-	38	29	34.8	44.4
P Avg	51.2	42.5	21.8	31.7	58	43.5	30.8	60.2	81.7	44.4	43.3	53.1	55.9	65.6	47.6	25.2	96.5	42.8	80	39.5	86.4	44.8	37.9	45	49.6
P+1	65.5	59.7	51.8	53.2	67.3	68	48	80.6	89.7	59.3	68.5	64.7	62.4	62.2	64.9	39	96.6	55.7	83.9	51.8	87.4	58	69.5	64.3	64.4
P+2	44.5	-	38.8	-	-	43.6	-	55.3	-	49.3	-	-	-	-	-	32.6	-	48.2	-	41.9	-	49.6	41.1	-	-
P+3	42.5	-	30.3	41.4	-	39.2	41.6	50.1	80.7	32.6	38.4	52.4	-	-	34.1	25.3	-	48.5	-	36.4	-	40.5	33.9	46.7	49.8
P+ Avg	58.1	59.7	40.3	47.3	67.3	50.3	44.8	62	85.2	47.1	53.5	58.6	62.4	62.2	49.5	32.3	96.6	50.8	83.9	43.4	87.4	49.4	**48.2**	55.5	57.1
S1	60.4	57.2	55.5	54.5	70.6	67.4	33.3	70.4	90.6	52.6	46.2	59.8	63.9	64.9	37.6	30.2	97	49.2	83.6	50.4	75.6	61.9	50	62.9	63.8
S2	41.7	-	40.8	-	-	39.6	-	59	-	48.1	-	-	-	-	-	24.9	-	47.6	-	34.8	-	46	34.5	-	-
S3	37	-	36.2	32.2	-	30	24.8	50	80.1	30.5	37.2	44.1	-	-	22.2	19.6	-	43.9	-	39.1	-	44.6	20.1	42.4	32.4
S Avg	53.6	57.2	44.2	43.4	70.6	45.7	29.1	59.8	85.4	43.7	41.7	52	63.9	64.9	29.9	24.9	97	46.9	83.6	41.4	75.6	50.8	34.9	52.7	48.1
C1	64.3	66.5	55.8	49.7	61.7	69.6	42.7	82.4	92.2	63.3	64.1	68.7	72.3	70.6	62.6	21.3	97	58.7	86.5	55.2	92.4	61.4	17.3	66.8	63.4
C2	**46.5**	-	42.6	-	-	47.4	-	65.1	-	49.4	-	-	-	-	-	22.9	-	62.2	-	42.6	-	57.2	29.1	-	-
C3	**46.4**	-	41.9	41.8	-	43.9	36.3	58.7	82.5	37.8	48.9	60.5	-	-	34.1	20.1	-	58.2	-	42.9	-	49.4	21.3	53.1	58.9
C Avg	59.8	66.5	**46.8**	45.8	61.7	**53.6**	39.5	**68.7**	87.4	**50.2**	56.5	**64.6**	**72.3**	70.6	48.4	21.4	97	**59.7**	**86.5**	46.9	**92.4**	56	22.6	**60**	**61.2**
Ours1	**67.3**	88.3	56.7	42.4	74	45.8	71.3	80.6	91.8	60.1	67.3	65.8	67.4	70.7	60.7	49.3	97	49.9	86.1	49.7	90.2	55.8	72.4	58.9	63.6
Ours2	45.4	-	35.4	-	-	48	-	63.8	-	42	-	-	-	-	-	40	-	49.5	-	42.6	-	45	42	-	-
Ours3	43.6	-	30.3	58.5	-	41.3	16	50.7	83.1	45.4	46.5	50.8	-	-	48.5	27.2	-	64	-	38.4	-	40.2	16	37.7	46.1
Ours Avg	60.6	**88.3**	40.8	**50.5**	74	43.5	**45.1**	65	**87.5**	49.2	**56.9**	58.3	67.4	**70.7**	**54.6**	**38.8**	97	54.5	86.1	43.6	90.2	47	43.5	48.3	54.8

Table 3: Coarse-to-fine grained Semantic Segmentation results on PartNet dataset (part-category mIoU %). The five methods are P:PointNet [38], P+:PointNet++ [40], S:SpiderCNN [60], C:PointCNN [25] and ours:MARNet. These methods are compared on three levels of segmentation 1:coarse-, 2:middle- and 3:fine-grained for 24 different categories. The average across the three levels and the average across the shape categories reveal that MARNet performs better or on-par with other methods. (best values are highlighted).
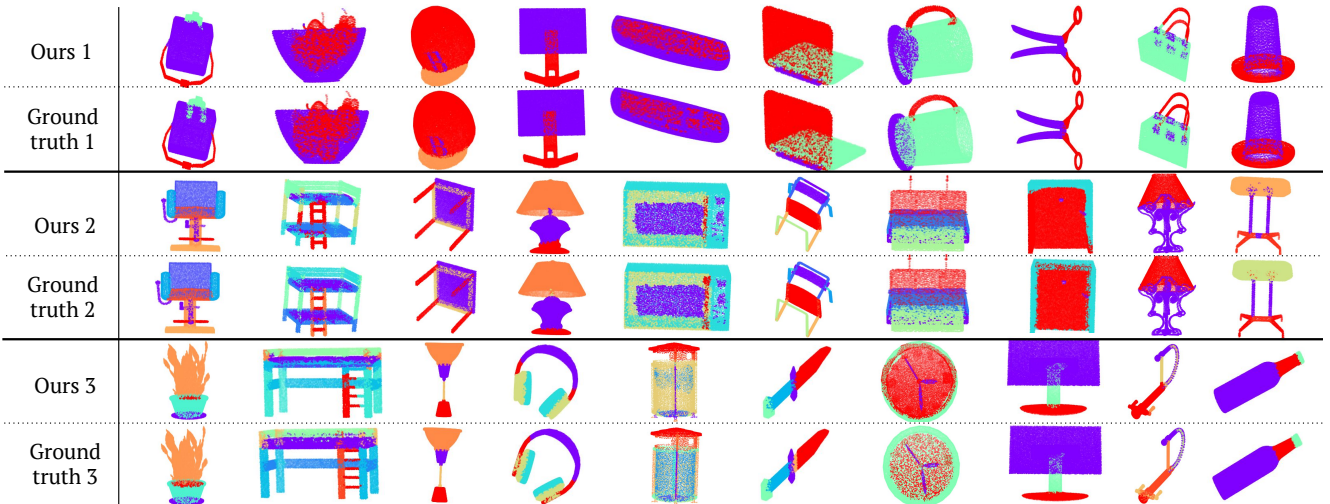

Figure 5: Qualitative results on PartNet Coarse-to-fine grained semantic segmentation. Top, Middle, and Bottom rows show MARNet prediction vs ground truths of 1:coarse-, 2:middle-, and 3:fine-grained segmentation respectively.

Specifically, we conduct experiments to ablate on 1) various components in the network, 2) the number of input points to the network, and 3) noi2y input data. Further, we outline the advantages of MARNet in terms of memory and computational complexity (measured by the number of parameters and FLOPs, respectively). All the ablation studies are conducted on ModelNet40 [58] classification dataset.

5.1. Componentwise contributions in MARNet

In this section, we conduct experiments to dissect the contributions of different components in MARNet and show the results in Tab. 4. First, we train only the backbone network (modified PointNet++, as in Sec. 3.2) without data augmentation (DA) and observe the performance as 89.1%. Adding data-augmentation (Sec. 4.1) along with the Backbone model (Backbone+DA) improves the performance by +1.1%. Next, we add FCR, FRE stages to formulate MARNet, but without residual function (MARNet(w/o $r(\cdot)$)) model. This simple architecture add-on significantly improves the classification accuracy by +2.4%. Such performance increase asserts that allowing multi-level features to interact with cross-referencing and re-encoding layers provides more contextual information. Next, we introduce residual functions $r(\cdot)$ in all stages of MARNet (as explained in Sec. 3.2, 3.3, 3.4) to improve the training stability and avoid vanishing/exploding gradients. MARNet(w/ $r(\cdot)$) variant passes unmodified gradients from final classification

Model	DA	point-wise trans. $\phi(\cdot)$	residual trans. $r(\cdot)$	voting	OA(%)
Backbone-only					89.1
Backbone+DA	✓				90.2
MARNet(w/o $r(\cdot)$)	✓	✓			92.6
MARNet(w/ $r(\cdot)$)	✓	✓	✓		93.4
MARNet	✓	✓	✓	✓	**93.9**

Table 4: Ablation studies on various components of MARNet. Backbone-only = modified PointNet++, Backbone+DA = Baseline (modified PointNet++ backbone as in Sec. 3.2) with Data Augmentation (DA), MARNet(w/o $r(\cdot)$) = MARNet without residual connections in all three stages, MARNet(w/ $r(\cdot)$) = including the residual connections and MARNet = complete MARNet evaluated with voting evaluation, *OA* = Overall Accuracy
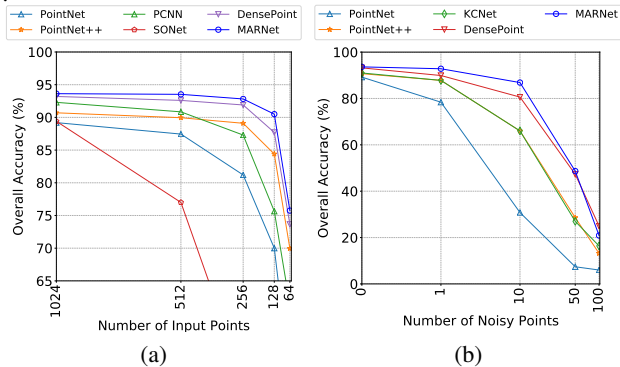
.

Figure 6: Performance comparison of various models on different number of input points (Left) and Noisy input points (Right). MARNet exhibits only minor degradation in performance compared to other state-of-the-art models.

layer to all levels in 3 stages and improves the performance further by +0.8%. Incorporating a voting mechanism for testing (as followed in [27, 38]) adds +0.5% improvement to MARNet.

5.2. Ablation on number of input points to the model

The resolution of point clouds tends to affect the performance of the models. Typically, high-resolution point clouds may provide richer and unambiguous features to aid classification. To study the effect of the different number of input points, we test MARNet with varying number of input points selected by the farthest point sampling method. Specifically, we test MARNet on 512, 256, 128 and 64 input points and show the results in Fig. 6a. Compared to several state-of-the-art methods [38, 40, 31, 23, 55, 27], MARNet exhibits robust performance with better classification accuracy than other methods even for sparse input points. For example, even with 128 points, MARNet's accuracy drops only by 2.9 % from the overall accuracy. We hypothesize that the ability to preserve the features at all levels of abstraction helps MARNet achieve such robustness.

Model	#params	#FLOPs	OA
PointNet [38]	3.50M	440M	89.2
KCNet [44]	0.9M	-	91.0
PointNet++ [40]	1.48M	1684M	91.9
SpecGCN [53]	2.05M	1112M	92.1
DGCNN [55]	1.84M	2767M	92.2
PointCNN [25]	**0.60M**	1581M	92.2
PCNN [31]	8.20M	294M	92.3
DensePoint [27]	0.67M	651M	93.2
MARNet	1.13M	1040M	**93.9**
MARNet Lite	0.68M	**260M**	93.2

Table 5: Comparison of different models in terms of the number of parameters and computational complexity (FLOPs). Here, *OA* = Overall Accuracy

5.3. Ablation on noisy input points

3D point clouds acquired by real sensors could possess noisy points, unlike computer-generated CAD models. To analyze the robustness towards the noise, we simulate noisy points by including points sampled from a random uniform distribution between -1 and 1. We test MARNet and other state-of-the-art methods by including 1, 10, 50, and 100 noisy points to the actual 1024 input points. Fig. 6b reveals that MARNet performs on-par to DensePoint and better than other state-of-the-art methods.

5.4. Model Complexity

Tab. 5 compares the complexities of different models in terms of the number of parameters and floating-point operations (FLOPs). MARNet achieves 93.9% accuracy and has a competitive model complexity with 1.13M parameters and ∼1 GFLOPs. A significant number of parameters are contributed from our Backbone's (PointNet++ [40]) MSG module. To avoid this, we propose a variant of MARNet, namely "MARNet Lite", by replacing the MSG layers with a single scale grouping layer. The resulting MARNet Lite has significantly less number of FLOPs and parameters with only a negligible drop in the accuracy.

6. Conclusion

In this work, we proposed a novel three-stage deep learning architecture named MARNet to extract and refine the multi-level abstract features for point-cloud analysis. Not only MARNet aggregates features of different granularity, but it also preserves earlier layer features till the final layer and propagates unmodified gradients to earlier layers. With these unique advantages, MARNet shows promising improvements over state-of-the-arts in the tasks of 3D shape classification and semantic segmentation. As the design of MARNet architecture is generic, we intend to study the applicability of this in other 3D point cloud tasks such as point cloud registration and object detection in our future work.

References

[1] Laurent Besacier, Etienne Barnard, Alexey Karpov, and Tanja Schultz. Automatic speech recognition for under-resourced languages: A survey. *Speech Communication*, 56:85–100, 2014. 1

[2] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. 6

[3] Chao Chen, Guanbin Li, Ruijia Xu, Tianshui Chen, Meng Wang, and Liang Lin. ClusterNet: Deep hierarchical cluster network with rigorously rotation-invariant representation for point cloud analysis. In *CVPR*, pages 4994–5002, 2019. 3

[4] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4D spatio-temporal convnets: Minkowski convolutional neural networks. *arXiv preprint arXiv:1904.08755*, 2019. 1, 2

[5] Yueqi Duan, Yu Zheng, Jiwen Lu, Jie Zhou, and Qi Tian. Structural relational reasoning of point clouds. In *CVPR*, pages 949–958, 2019. 2

[6] Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural networks. In *AAAI*, volume 33, pages 3558–3565, 2019. 3

[7] Yifan Feng, Zizhao Zhang, Xibin Zhao, Rongrong Ji, and Yue Gao. GVCNN: Group-view convolutional neural networks for 3D shape recognition. In *CVPR*, pages 264–272, 2018. 2

[8] Matheus Gadelha, Rui Wang, and Subhransu Maji. Multiresolution tree networks for 3D point cloud processing. In *ECCV*, pages 105–122, 2018. 2

[9] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving. In *CVPR*, pages 3354–3361, 2012. 1

[10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 1, 4

[11] Pedro Hermosilla, Tobias Ritschel, Pere-Pau Vázquez, Alvar Vinacua, and Timo Ropinski. Monte carlo convolution for learning on non-uniformly sampled point clouds. *ACM Trans. Graph.*, 37(6):235:1–235:12, 2018. 2

[12] Sepp Hochreiter. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116, 1998. 5

[13] Binh-Son Hua, Minh-Khoi Tran, and Sai-Kit Yeung. Pointwise convolutional neural networks. In *CVPR*, 2018. 2

[14] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, pages 2261–2269, 2017. 1, 2, 6

[15] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, page arXiv:1412.6980, Dec. 2014. 5

[16] Greg Kipper and Joseph Rampolla. *Augmented Reality: an emerging technologies guide to AR*. Elsevier, 2012. 1

[17] Roman Klokov and Victor Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3D point cloud models. In *ICCV*, pages 863–872, 2017. 2, 6

[18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet classification with deep convolutional neural networks. In *NeurIPS*, pages 1106–1114, 2012. 1, 2, 4, 11

[19] Shiyi Lan, Ruichi Yu, Gang Yu, and Larry S Davis. Modeling local geometric structure of 3D point clouds using Geo-CNN. *arXiv preprint arXiv:1811.07782*, 2018. 2, 3, 6

[20] Truc Le and Ye Duan. Pointgrid: A deep network for 3d shape understanding. In *CVPR*, pages 9204–9214, 2018. 1, 2

[21] Haun Lei, Naveed Akhtar, and Ajmal Mian. Spherical convolutional neural network for 3d point clouds. 2018. 3

[22] Huan Lei, Naveed Akhtar, and Ajmal Mian. Octree guided cnn with spherical kernels for 3D point clouds. *arXiv preprint arXiv:1903.00343*, 2019. 2

[23] Jiaxin Li, Ben M. Chen, and Gim Hee Lee. SO-Net: Self-organizing network for point cloud analysis. In *CVPR*, pages 9397–9406, 2018. 6, 8

[24] Ruoyu Li, Sheng Wang, Feiyun Zhu, and Junzhou Huang. Adaptive graph convolutional neural networks. In *AAAI*, 2018. 3

[25] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. PointCNN: Convolution on x-transformed points. In *NeurIPS*, pages 820–830, 2018. 1, 3, 6, 7, 8

[26] Jinxian Liu, Bingbing Ni, Caiyuan Li, Jiancheng Yang, and Qi Tian. Dynamic points agglomeration for hierarchical point sets learning. In *ICCV*, pages 7546–7555, 2019. 3

[27] Yongcheng Liu, Bin Fan, Gaofeng Meng, Jiwen Lu, Shiming Xiang, and Chunhong Pan. DensePoint: Learning densely contextual representation for efficient point cloud processing. In *ICCV*, pages 5239–5248, 2019. 1, 2, 3, 4, 6, 8

[28] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *CVPR*, pages 8895–8904, 2019. 3

[29] Chao Ma, Yulan Guo, Jungang Yang, and Wei An. Learning multi-view representation with LSTM for 3D shape recognition and retrieval. *IEEE TMM*, 2018. 2

[30] Jiageng Mao, Xiaogang Wang, and Hongsheng Li. Interpolated convolutional networks for 3d point cloud understanding. In *ICCV*, pages 1578–1587, 2019. 3

[31] Atzmon Matan, Maron Haggai, and Lipman Yaron. Point convolutional neural networks by extension operators. *ACM TOG*, 37(4):1–12, 2018. 2, 3, 6, 8

[32] Daniel Maturana and Sebastian Scherer. VoxNet: A 3D convolutional neural network for real-time object recognition. In *IROS*, pages 922–928, 2015. 1, 2

[33] Kaichun Mo, Shilin Zhu, Angel X Chang, Li Yi, Subarna Tripathi, Leonidas J Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *CVPR*, pages 909–918, 2019. 6

[34] Daniel W Otter, Julian R Medina, and Jugal K Kalita. A survey of the usages of deep learning in natural language processing. *arXiv preprint arXiv:1807.10854*, 2018. 1

[35] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017. 5

[36] Javier Andreu Perez, Fani Deligianni, Daniele Ravi, and Guang-Zhong Yang. Artificial intelligence and robotics. *arXiv preprint arXiv:1803.10813*, 2018. 1

[37] Adrien Poulenard, Marie-Julie Rakotosaona, Yann Ponty, and Maks Ovsjanikov. Effective rotation-invariant point CNN with spherical harmonics kernels. In *3DV*, pages 47–56, 2019. 3

[38] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, pages 77–85, 2016. 1, 2, 4, 6, 7, 8, 12

[39] 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 *CVPR*, pages 5648–5656, 2016. 1, 2

[40] Charles R. Qi, Li Yi, Hao su, and Leonidas J. Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, pages 5099–5108, 2017. 1, 2, 3, 4, 5, 6, 7, 8, 12

[41] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. OctNet: Learning deep 3D representations at high resolutions. In *CVPR*, pages 6620–6629, 2017. 2

[42] 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. 4

[43] Radu Alexandru Rosu, Peer Schütt, Jan Quenzel, and Sven Behnke. Latticenet: Fast point cloud segmentation using permutohedral lattices. *arXiv preprint arXiv:1912.05905*, 2019. 2

[44] Yiru Shen, Chen Feng, Yaoqing Yang, and Dong Tian. Mining point cloud local structures by kernel correlation and graph pooling. In *CVPR*, pages 4548–4557, 2018. 3, 6, 8

[45] M. Simonovsky and N. Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *CVPR*, pages 29–38, 2017. 3, 6

[46] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, pages 1–14, 2015. 1

[47] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3D shape recognition. In *ICCV*, pages 945–953, 2015. 1, 2

[48] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020. 1

[49] Shuyang Sun, Jiangmiao Pang, Jianping Shi, Shuai Yi, and Wanli Ouyang. Fishnet: A versatile backbone for image, region, and pixel level prediction. In *Advances in Neural Information Processing Systems*, pages 760–770, 2018. 1, 2, 4

[50] Gusi Te, Wei Hu, Amin Zheng, and Zongming Guo. RGCNN: Regularized graph CNN for point cloud segmentation. In *ACM MM*, pages 746–754, 2018. 3

[51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. 1

[52] Chu Wang, Marcello Pelillo, and Kaleem Siddiqi. Dominant set clustering and pooling for multi-view 3D object recognition. *arXiv preprint arXiv:1906.01592*, 2019. 2

[53] Chu Wang, Babak Samari, and Kaleem Siddiqi. Local spectral graph convolution for point set feature learning. In *ECCV*, pages 52–66, 2018. 3, 6, 8

[54] 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 TOG*, 36(4):72, 2017. 2

[55] 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 Trans. Graph.*, pages 1–13, 2019. 3, 6, 8

[56] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. *arXiv preprint arXiv:1811.07246*, 2018. 2

[57] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D shapenets: A deep representation for volumetric shapes. In *Proceedings of CVPR*, pages 1912–1920, 2015. 1, 2

[58] 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*, pages 1912–1920, 2015. 5, 7

[59] Saining Xie, Sainan Liu, Zeyu Chen, and Zhuowen Tu. Attentional shapecontextnet for point cloud recognition. In *CVPR*, pages 4606–4615, 2018. 3, 6

[60] Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. SpiderCNN: Deep learning on point sets with parameterized convolutional filters. In *ECCV*, pages 87–102, 2018. 2, 6, 7

[61] Jiancheng Yang, Qiang Zhang, Bingbing Ni, Linguo Li, Jinxian Liu, Mengdie Zhou, and Qi Tian. Modeling point clouds with self-attention and gumbel subset sampling. *arXiv preprint arXiv:1904.03375*, 2019. 2

[62] Ze Yang and Liwei Wang. Learning relationships for multi-view 3D object recognition. In *ICCV*, pages 7505–7514, 2019. 1, 2

[63] Tan Yu, Jingjing Meng, and Junsong Yuan. Multi-view harmonized bilinear network for 3D object recognition. In *CVPR*, pages 186–194, 2018. 1, 2

[64] Yiming Zeng, Yu Hu, Shice Liu, Jing Ye, Yinhe Han, Xiaowei Li, and Ninghui Sun. RT3D: Real-time 3D vehicle detection in lidar point cloud for autonomous driving. *IEEE Robotics and Automation Letters*, 3(4):3434–3440, 2018. 1

[65] Kuangen Zhang, Ming Hao, Jing Wang, Clarence W. de Silva, and Chenglong Fu. Linked dynamic graph CNN: Learning on point cloud via linking hierarchical features. *arXiv preprint arXiv:1904.10014*, 2019. 3

10

Supplementary Material

1. Outline

In the supplementary section, we report additional ablation studies of our model on the following aspects:

- varying number of groups (N_g) in Grouped convolution based MLP $\phi(\cdot)$ (Sec. 2.1)

- varying number of levels in the backbone, Feature Cross-Referencing (FCR) and the Feature Re-Encoding (FRE) stages (Sec. 2.2)

- Inference time and memory requirements (Sec. 2.3)

More qualitative results on PartNet part-segmentation can be found in Fig. 1. Additionally, we describe the exact network architecture of MARNet and its training details in Sec. 3, which will be of help to reproduce the work.

2. Ablation studies (contd.,)

2.1. Varying number of groups (N_g) in Grouped Convolution based MLPs $\phi(\cdot)$

In grouped convolution [18], input channels are divided into N_g groups and convolution kernel is applied separately on them. Then, the output features are generated independently and concatenated to obtain the final output. This type of grouping imposes sparsity in the connections between the nodes and helps to reduce the model complexity in terms of number of parameters.

Tab. 1 shows the influence of different N_g on model complexity (#parameters, #floating point operations (FLOPs)) and accuracy of our model. We experiment with values of N_g between 1 and 8. As N_g increases, #parameters and #FLOPs decreases due to the sparsification of the connections between the nodes. We obtained the highest accuracy for $N_g = 2$ and thus, we select this configuration as our optimal model.

N_g	#parameters	#FLOPs	O.A. (%)
1	1.85M	2080M	93.2
2	1.13M	1040M	**93.9**
4	0.84M	580M	93.4
8	0.67M	380M	92.5

Table 1: MARNet complexity and performance for N_g groups in grouped convolution. Here, *OA* = Overall Accuracy.

2.2. Varying number of levels in MARNet

We report the model complexity (#parameters, #FLOPs) and its performance for different number of levels in MARNet. As mentioned in Sec. 3 of the main paper, MARNet has l levels of abstraction (down-sampling) in the Backbone

#levels	#parameters	#FLOPs	O.A. (%)
BB = 3, FCR = 2, FRE = 2	0.80M	1320M	93.1
BB = 4, FCR = 3, FRE = 3	1.13M	1040M	**93.9**
BB = 5, FCR = 4, FRE = 4	2.06M	3060M	93.5
BB = 6, FCR = 5, FRE = 5	3.42M	8020M	93.0

Table 2: MARNet complexity and performance for varying number of #levels in MARNet. Here, *BB, FCR, GRE* denote number of Backbone-, FCR- and FRE-levels respectively. The number of FC layers = 3 in all cases. *OA* = Overall Accuracy.

model	time	memory	O.A. (%)
MARNet	24ms	1951MB	93.9
MARNet Lite	8ms	1055MB	93.2

Table 3: Inference time and GPU memory requirements for MARNet and MARNet Lite. Batch size of 32 is used. "time" denotes average inference time per test sample.

stage given by $\{L_{bb}^0, L_{bb}^1, \ldots, L_{bb}^{l-1}\}$. For every level in the Backbone (except L_{bb}^0), one corresponding level is added in FCR and FRE stages respectively.

Tab. 2 gives the performance of MARNet by varying number of levels. MARNet with 10 levels (Backbone = 3, FCR = 2, FRE = 2, FC = 3) lacks sufficient capacity to learn multi-abstraction information and therefore, it underperforms. MARNet with 13 levels (Backbone = 4, FCR = 3, FRE = 3, FC layers = 3) gives the best accuracy of 93.9%. Further increasing the levels increases the model complexity and overfits to the training data while giving sub-par performance.

2.3. Inference time and memory requirements

We compare the computational requirements of our optimal models (MARNet and MARNet Lite with 13 levels: Backbone = 4, FCR = 3, FRE = 3, FC layers = 3) in Tab. 3. The test is executed on a Nvidia GTX 1080Ti GPU with batch size of 32.

We observe that MARNet is competitive in terms of inference time and memory requirement. However, MARNet Lite is 3 times faster than MARNet and achieves a near real-time inference speed with a little drop in accuracy.

3. Network Architecture details

In this section, we provide the exact settings used for training MARNet for different tasks to facilitate reproducibility. In particular, we give layer-by-layer configurations for the following :

1. MARNet for classification task (Tab. 4)

2. MARNet for part-segmentation task (Tab. 5)

3. MARNet lite for classification task (Tab. 6)

Firstly, we explain the layer types used in our network. A brief description of the network components along with the hyper-parameters are explained below.

PointNetSetAbstraction:
hyper-params: (#input channels, grouping radius, #samples, [MLP output dims])
We use Ball Point Query to select the points within the grouping radius and choose #samples points among them to group and aggregate locally using PointNet [38]. Point-Net is implemented as series of MLPs with output dimensions given in *MLP output dims* list. Max pooling is used to aggregate the local point features.

PointNetSetAbstractionMsg:
hyper-params: (#input channels, [list of grouping radii], [list of #samples], [[list of MLP output dims]])
This layer is Multi-Scale Grouping (MSG) version of the *PointNetSetAbstraction* layer. The grouping, transformation and aggregation operations of PointNet are applied for multiple radii in parallel and their features are concatenated in the end.

FeatureCrossReferencing:
hyper-params: (#input channels, [MLP output dims])
The backbone features are concatenated with the input of this layer. The features are passed through a set of MLPs and a reduction function before upsampling the points. To project the point features on the upsampled points using 3-nearest neighbor (3NN) interpolation technique. The 3NN technique interpolates the points according to the inverse of the weighted distance between the points.

FeatureReEncoding:
hyper-params: (#input channels, grouping radius, #samples, [MLP output dims])
The features from backbone and FCR layers are concatenated to this layer's input. Same set of sampled points from the backbone are passed to this layer. The grouping and aggregation operation is similar to the *PointNetSetAbstraction* layer.

PointNetFeaturePropagation:
hyper-params: (#input channels, [MLP output dims])
This layer is only employed in our part-segmentation network. It is the feature propagation layer of PointNet++ [40]. The features from the encoder at the same level are concatenated, transformed and upsampled. The transformation is implemented as a series of MLPs.

FullyConnectedLayer:
hyper-params: (#input channels, #output channels,

dropout ratio %)
The FC layer conprises of a linear layer followed by a batch-normalization layer and a ReLU non-linearity. Dropout is applied after every layer except for the final layer which is used to make the predictions.

All the point-wise transformation layers in the Backbone, FCR and FRE stages use grouped convolutions with number of groups $N_g = 2$ throughout.

In the backbone stage, apart from the input channels from the previous layer, 6 more channels (xyz+normals) are concatenated to features before applying PointNet [38] locally.

3.1. MARNet Classification Network

Our ModelNet40 classification network consists of 13 levels in total: #Backbone levels = 4, #FCR levels = 3, #FRE levels = 3, number of fully connected (FC) layers = 3. The details of our network design are mentioned in Tab. 4.

3.2. MARNet Part-Segmentation Network

Tab. 5 outlines our network design for the PartNet part-segmentation task. It consists of 16 levels: #Backbone levels = 4, #FCR levels = 3, #FRE levels = 3, PointNetFeature-Propagation layers = 4, and #FC layers = 2.

3.3. MARNet Lite for classification task

This is the lite version of MARNet in terms of #parameters and #FLOPS. It comprises of the same layout as MARNet classification network except that the MSG layers in the backbone are replaced by the SSG (single scale grouping) layers. MARNet Lite's configuration is listed in Tab. 6.
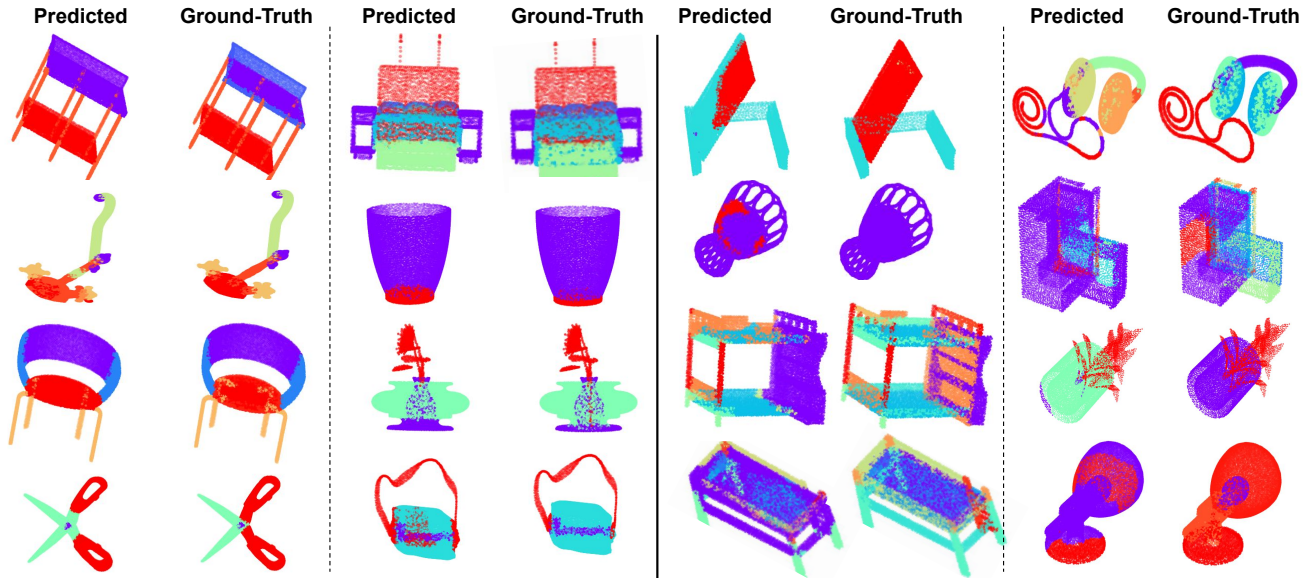
Figure 1: Additional qualitative results of MARNet on PartNet dataset. Columns 1 and 2 show the accurate predictions of MARNet with respect to the ground-truth. Columns 3 and 4 show some of the incorrect predictions.

Layer	Layer Type	Layer Parameters	Output	S.C.
		Backbone Stage		
BB1	PointNetSetAbstractionMsg	$0, [0.1, 0.2, 0.4], [16, 32, 128], [[16, 16, 16], [16, 16, 16], [32, 32, 32]]$	$(64, 512)$	
BB2	PointNetSetAbstractionMsg	$64, [0.2, 0.4, 0.6], [32, 64, 128], [[32, 32, 32], [32, 32, 32], [64, 64, 64]]$	$(128, 128)$	
BB3	PointNetSetAbstractionMsg	$128, [0.6, 0.8, 0.9], [64, 96, 128], [[64, 64, 64], [64, 64, 64], [128, 128, 128]]$	$(256, 32)$	
BB4	PointNetSetAbstraction	$256, None, None, [256]$	$(256, 1)$	
		Feature Cross Referencing Stage		
FCR1	FeatureCrossReferencing	$256, [128, 128]$	$(128, 32)$	
FCR2	FeatureCrossReferencing	$128 + 256, [64, 64]$	$(64, 128)$	BB3
FCR3	FeatureCrossReferencing	$64 + 128, [32, 32]$	$(32, 512)$	BB2
		Feature Re-Encoding Stage		
FRE1	FeatureReEncoding	$32 + 64, 0.4, 32, [96, 96]$	$(96, 128)$	BB1
FRE2	FeatureReEncoding	$96 + 64 + 128, 0.8, 32, [288, 288]$	$(288, 32)$	BB2, FCR2
FRE3	FeatureReEncoding	$288 + 128 + 256, None, None, [672, 672]$	$(672, 1)$	BB3, FCR1
		Classification Stage		
FC1	FullyConnectedLayer	$672, 512, 0.4$	$(512, 1)$	
FC2	FullyConnectedLayer	$512, 256, 0.5$	$(256, 1)$	
FC3	FullyConnectedLayer	$256, \#classes$	$(\#classes, 1)$	

Table 4: Network Architecture of MARNet Classification Network. Here, *Output* = (#output channels, #output points), *S.C.* = shortcut connections indicating features passed from previous layers.

Layer	Layer Type	Layer Parameters	Output	S.C.
		Backbone Stage		
BB1	PointNetSetAbstractionMsg	$0, [0.1, 0.2, 0.4], [16, 32, 128], [[16, 16, 16], [16, 16, 16], [32, 32, 32]]$	$(64, 512)$	
BB2	PointNetSetAbstractionMsg	$64, [0.2, 0.4, 0.6], [32, 64, 128], [[32, 32, 32], [32, 32, 32], [64, 64, 64]]$	$(128, 128)$	
BB3	PointNetSetAbstractionMsg	$128, [0.6, 0.8, 0.9], [64, 96, 128], [[64, 64, 64], [64, 64, 64], [128, 128, 128]]$	$(256, 32)$	
BB4	PointNetSetAbstraction	$256, None, None, [256]$	$(256, 1)$	
		Feature Cross Referencing Stage		
FCR1	FeatureCrossReferencing	$256, [128, 128]$	$(128, 32)$	
FCR2	FeatureCrossReferencing	$128 + 256, [64, 64]$	$(64, 128)$	BB3
FCR3	FeatureCrossReferencing	$64 + 128, [32, 32]$	$(32, 512)$	BB2
		Feature Re-Encoding Stage		
FRE1	FeatureReEncoding	$32 + 64, 0.4, 32, [96, 96]$	$(96, 128)$	BB1
FRE2	FeatureReEncoding	$96 + 64 + 128, 0.8, 32, [288, 288]$	$(288, 32)$	BB2, FCR2
FRE3	FeatureReEncoding	$288 + 128 + 256, None, None, [672, 672]$	$(672, 1)$	BB3, FCR1
		Segmentation Stage		
FP1	PointNetFeaturePropagation	$672 + 288, [256, 256]$	$(256, 32)$	FRE2
FP2	PointNetFeaturePropagation	$256 + 96, [256, 128]$	$(128, 128)$	FRE1
FP3	PointNetFeaturePropagation	$128 + 32, [128, 128]$	$(128, 512)$	FCR3
FP4	PointNetFeaturePropagation	$128 + 6, [128, 128]$	$(128, 1024)$	
		Point-wise Classification Stage		
FC1	FullyConnectedLayer	$128, 128, 0.5$	$(128, 1024)$	
FC2	FullyConnectedLayer	$128, \#parts$	$(\#parts, 1024)$	

Table 5: Network Architecture of MARNet Part-Segmentation Network. Here, *Output* = (#output channels, #output points), *S.C.* = shortcut connections indicating features passed from previous layers.

Layer	Layer Type	Layer Parameters	Output	S.C.
		Backbone Stage		
BB1	PointNetSetAbstraction	$0, 0.2, 32, [32, 32, 32]$	$(64, 512)$	
BB2	PointNetSetAbstraction	$32, 0.4, 32, [64, 64, 64]$	$(128, 128)$	
BB3	PointNetSetAbstraction	$64, 0.8, 32, [128, 128, 128]$	$(256, 32)$	
BB4	PointNetSetAbstraction	$128, None, None, [256]$	$(256, 1)$	
		Feature Cross Referencing Stage		
FCR1	FeatureCrossReferencing	$256, [128, 128]$	$(128, 32)$	
FCR2	FeatureCrossReferencing	$128 + 128, [64, 64]$	$(64, 128)$	BB3
FCR3	FeatureCrossReferencing	$64 + 64, [32, 32]$	$(32, 512)$	BB2
		Feature Re-Encoding Stage		
FRE1	FeatureReEncoding	$32 + 32, 0.4, 32, [64, 64]$	$(96, 128)$	BB1
FRE2	FeatureReEncoding	$64 + 64 + 64, 0.8, 32, [192, 192]$	$(288, 32)$	BB2, FCR2
FRE3	FeatureReEncoding	$192 + 128 + 128, None, None, [448, 448]$	$(672, 1)$	BB3, FCR1
		Classification Stage		
FC1	FullyConnectedLayer	$448, 512, 0.4$	$(512, 1)$	
FC2	FullyConnectedLayer	$512, 256, 0.5$	$(256, 1)$	
FC3	FullyConnectedLayer	$256, \#classes$	$(\#classes, 1)$	

Table 6: Network Architecture of MARNet Lite Classification Network. Here, *Output* = (#output channels, #output points), *S.C.* = shortcut connections indicating features passed from previous layers.