

ViPFormer: Efficient Vision-and-Pointcloud Transformer for Unsupervised Pointcloud Understanding

Hongyu Sun, Yongcai Wang, Xudong Cai, Xuewei Bai and Deying Li

Abstract—Recently, a growing number of work design unsupervised paradigms for point cloud processing to alleviate the limitation of expensive manual annotation and poor transferability of supervised methods. Among them, CrossPoint follows the contrastive learning framework and exploits image and point cloud data for unsupervised point cloud understanding. Although the promising performance is presented, the unbalanced architecture makes it unnecessarily complex and inefficient. For example, the image branch in CrossPoint is $\sim 8.3x$ heavier than the point cloud branch leading to higher complexity and latency. To address this problem, in this paper, we propose a lightweight Vision-and-Pointcloud Transformer (ViPFormer) to unify image and point cloud processing in a single architecture. ViPFormer learns in an unsupervised manner by optimizing intra-modal and cross-modal contrastive objectives. Then the pretrained model is transferred to various downstream tasks, including 3D shape classification and semantic segmentation. Experiments on different datasets show ViPFormer surpasses previous state-of-the-art unsupervised methods with higher accuracy, lower model complexity and runtime latency. Finally, the effectiveness of each component in ViPFormer is validated by extensive ablation studies. The implementation of the proposed method is available at <https://github.com/auniquesun/ViPFormer>.

I. INTRODUCTION

Point cloud understanding is a crucial problem which has attracted widespread attention for its values in autonomous driving, mixed reality, and robotics. There are three common tasks in point cloud understanding: 3D object classification [1], [2], semantic segmentation [3], [4], [5] and object detection [6], [7], [8], [9]. A large portion of previous methods design different neural networks and learn from large-scale *annotated* data for point cloud understanding tasks. However, point cloud labels are rare in most scenarios and acquiring them is time-consuming and expensive.

Hence, in recent years, researchers have begun to shift their attention to developing unsupervised methods for point cloud understanding, without the need of hand-crafted annotations. Unsupervised methods are designed in different ways, such as auto-encoders [10], mask auto-encoders [11], [12], reconstruction [13], occlusion completion [14], [15], GANs [16], [17], [18] and contrastive learning [19], [20], [21], [22], [23], etc.

Currently, a growing number of methods embrace contrastive learning because it is a simple yet effective frame-

All authors are with the Department of Computer Science, School of Information, Renmin University of China, Beijing 100872, China. Corresponding author: Yongcai Wang. {sunhongyu, ycw, xudongcai, bai_xuewei, deyingli}@ruc.edu.cn

This work was supported in part by the National Natural Science Foundation of China under Grants No. 61972404 and No. 12071478, and Public Computing Cloud, Renmin University of China.




Fig. 1: Comparison of classification accuracy and the number of parameters of different unsupervised methods on ModelNet40 (M) and ScanObjectNN (S), respectively.

work and has shown improvements in vision [24] and language processing [25]. This framework can be easily extended to incorporate multi-modal data to further exploit richer semantics in multi-modal data to improve performances. Inspired by the success of vision + language [26], [27] and video + audio [28], [29], point cloud understanding powered by cross-modal data has drawn research interests.

CrossPoint [22] takes images and point clouds as inputs and follows the contrastive framework for unsupervised point cloud processing. It utilizes ResNet50 [30] as image feature extractor and PointNet [31]/DGCNN [32] as point cloud feature extractor. [33] uses CNN and U-Net [34] architecture in the image branch and PointNet++ [35] architecture in the point cloud branch for contrastive learning. Although promising performances are obtained, the unbalanced image and point cloud processing architecture makes them unnecessarily complex and inefficient. For example, in [22], the point cloud branch has 3M parameters while the image branch has 25M. The image processing branch is $\sim 8.3x$ heavier than the point cloud one and consumes much more time.

The different and unbalanced architecture when dealing with data from different modalities is often neglected in academic studies but is a critical problem in practice because it severely hinders efficiency. However, it is possible to design a unified and efficient architecture to process cross-modal data since Transformer [36] has shown the flexibility and superiority in vision [37], [38], [24] and language [39], [40], [41] modeling. And recently, Point-BERT [11] and Point-MAE [42] show point clouds can be sampled to groups then processed by Transformer and the performances are promising.

In this paper, we propose an efficient **Vision-and-Pointcloud Transformer (ViPFormer)** for unsupervised point cloud understanding. ViPFormer unifies image and point

cloud processing in a single architecture, which ensures the two branches have the same size and complexity. Then it follows the contrastive learning framework to optimize image and point cloud feature representations. Finally, the learned representations are transferred to target tasks like 3D point cloud classification and semantic segmentation.

ViPFormer is evaluated on a wide range of point cloud understanding tasks, including synthetic and real-world 3D shape classification, object part segmentation. Results show it not only reduces the model complexity and running latency but also outperforms all existing unsupervised methods. The major contributions of this paper are summarized as follows:

- We propose ViPFormer, handling image and point cloud data in a unified architecture, simplifies the model complexity, reduces running latency and boosts overall performances for unsupervised point cloud understanding.
- We show that ViPFormer can be generalized better to different tasks by simultaneously optimizing intra-modal and cross-modal contrastive objectives.
- The proposed method is validated on various downstream tasks, e.g., it achieves 90.7% classification accuracy on ScanObjectNN, leading CrossPoint by 9%, and surpassing the previous best performing unsupervised method by 0.7% with $\sim 77\%$ fewer parameters. Similarly, ViPFormer reaches a 93.9% score on ModelNet40, outperforming the previous state-of-the-art method and reducing the number of parameters by 24%.
- We conduct extensive ablation studies to clarify the advantages of the architecture design, contrastive optimization objectives, and unsupervised learning strategy.

II. RELATED WORK

Unsupervised Point Cloud Understanding Unsupervised learning becomes more and more popular since it can unleash the potential of large-scale unlabeled data and save considerable costs. According to the pretext task, unsupervised methods for point cloud understanding can be classified into generative models and discriminative models. Generative models usually learn the latent representations of point clouds by predicting some parts or all of the input data. The assumption is that only after the model understands the point cloud can it predict the occluded parts or generate the entire point cloud. Auto-encoders like FoldingNet [10], GANs like LRG [17], URL [18] and 3D GAN [16], reconstruction methods like JigSaw [13], cTree [43], all of them generate a whole point cloud and maximize the similarity with the input point cloud. Mask encoders like OcCo [14], PointBERT [11], Point-MAE [12] complete the masked parts of a point cloud to keep it same as the input. On the other hand, the discriminative models aim to learn discriminative features from different object/semantic categories. Most of them follow the contrastive learning framework [19], [20], [21], [44], [23], [22], where CrossPoint [22] is the most relevant to us since it also fuses cross-modal data, images and point clouds, for point cloud understanding. However, the unbalanced feature extractors in CrossPoint caused much

higher running latency and model complexity. Instead, we propose Vision-and-Pointcloud Transformer to unify image and point cloud processing in a single architecture, reduce latency and boost performance.

Architecture for Image and Pointcloud Processing An image consists of regular and dense pixel grids, while a point cloud is a set of irregular, sparse and unordered points. The huge differences make it difficult to process images and point clouds in the same way. Researchers developed different architectures for image and point cloud processing. In many cases, CNNs are the first choices of image processing and PointNet [31] and its variants [35], [45] are good starts for point cloud processing. However, the situation has changed since the advent of Transformer [36]. Due to the notable improvements, Transformer quickly became the de facto standard architecture for language understanding tasks [39], [40], [41] then entered vision [46], [37], [38], [24] and 3D field [47], [48], [49], [50], [42]. Guo *et al.* and Zhao *et al.* proposed PCT [51] and Point Transformer [4], respectively, but their architectures were different from the standard Transformer [36] and can not be generalized to vision modality. Perceiver [52] and PerceiverIO [53] have taken important steps toward general architecture for various modalities (audio, image, point cloud). However, Perceiver and PerceiverIO learn in a supervised fashion. Differently, the proposed ViPFormer unifies image and point cloud processing in a single architecture and learns from large-scale unlabeled data.

III. METHODOLOGY

In this section, firstly, we introduce the overall architecture of ViPFormer. Secondly, we elaborate on its unsupervised learning strategy.

A. The Overall Architecture of ViPFormer

As Fig. 2 shows, ViPFormer consists of three components, which are a lightweight Input Adapter, Transformer Encoder and Output Adapter. In image and point cloud branches, modules with the same color are identical. And the images and point clouds are serialized in different ways.

Image and Point Cloud Preparing To exploit the power of Transformer [36] we need to convert images and point clouds into sequence data as Transformer requires. Inspired by ViT [37], we divide an image into small patches and then flatten them into a sequence. For example, an image I is of size $H \times W \times C_1$ and the patch size is Q , we can generate an image patch sequence $\mathbf{x}_i \in \mathbb{R}^{M \times (Q^2 \cdot C_1)}$, where $M = HW/Q^2$.

For a point cloud P of size $N \times C_2$, we convert it into a patch sequence as follows [11], [12]. First, the farthest point sampling (FPS) is applied to P to get G centers. Second, for each center, we search its k nearest neighbors (k NN) in P to aggregate local geometry information, resulting in a patch sequence $\mathbf{x}_p \in \mathbb{R}^{G \times (k \cdot C_2)}$.

Input Adapter We design a lightweight image patch adapter \mathbf{E}_I and a point patch adapter \mathbf{E}_P to project the sequences to high dimensional feature representations. \mathbf{E}_I

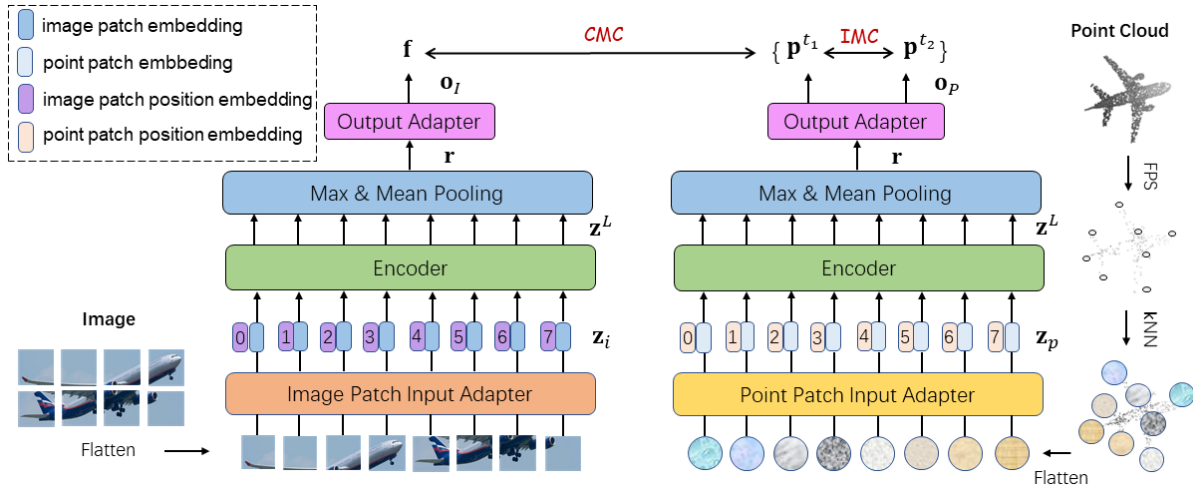


Fig. 2: The overall architecture of ViPFormer.

$\in \mathbb{R}^{(Q^2 \cdot C_1) \times D}$ is a linear layer and \mathbf{E}_P is a multi-layer perception (MLP). The outputs are denoted as image patch embeddings \mathbf{z}_i and point patch embeddings \mathbf{z}_p .

$$\mathbf{z}_i = \mathbf{x}_i \mathbf{E}_I, \quad \mathbf{z}_p = \mathbf{x}_p \mathbf{E}_P \quad (1)$$

Before being fed into Encoder, the position information is injected to \mathbf{z}_i and \mathbf{z}_p by adding the image patch position embeddings \mathbf{E}_I^{pos} and point patch position embeddings \mathbf{E}_P^{pos} .

$$\mathbf{z}_i = \mathbf{z}_i + \mathbf{E}_I^{pos}, \quad \mathbf{z}_p = \mathbf{z}_p + \mathbf{E}_P^{pos} \quad (2)$$

Encoder The image and point cloud branches share the Encoder architecture, which ensures image and point cloud processing have balanced complexity and low latency. ViPFormer Encoder consists of L stacked multi-head self-attention (MSA) and MLP layers. MLP has two layers with a GELU non-linear activation. LayerNorm (LN) is applied before MSA and MLP layers, while Dropout is applied after them.

$$\hat{\mathbf{z}}^l = \text{Dropout}(\text{MSA}(\text{LN}(\mathbf{z}^{l-1}))) + \mathbf{z}^{l-1}, \quad l = 1 \dots L \quad (3)$$

$$\mathbf{z}^l = \text{Dropout}(\text{MLP}(\text{LN}(\hat{\mathbf{z}}^l))) + \hat{\mathbf{z}}^l, \quad l = 1 \dots L \quad (4)$$

Before proceeding, the output sequence of ViPFormer Encoder needs to be converted into an object-level feature. We implement it by concatenating the max and mean value of \mathbf{z}^L .

$$\mathbf{r} = \text{Concat}(\text{Max}(\mathbf{z}^L), \text{Mean}(\mathbf{z}^L)) \quad (5)$$

Output Adapter The image and point cloud branches also share the Output Adapter. As suggested by SimCLR [54], a learnable nonlinear transformation between the encoder and the contrastive loss can improve the quality of feature representations. The Output Adapter is implemented by two consecutive Linear layers, preceding with BatchNorm (BN) and ReLU.

$$\hat{\mathbf{r}} = \text{Linear}(\text{ReLU}(\text{BN}(\mathbf{r}))) \quad (6)$$

$$\mathbf{o} = \text{Linear}(\text{ReLU}(\text{BN}(\hat{\mathbf{r}}))) \quad (7)$$

At this point, the input image I and point cloud P are transformed into image feature $\mathbf{f} = \mathbf{o}_I$ and point cloud feature $\mathbf{p} = \mathbf{o}_P$. We can use these features for unsupervised contrastive learning.

B. Unsupervised Contrastive Pretraining of ViPFormer

We conduct unsupervised pretraining for ViPFormer by introducing two contrastive objectives, intra-modal contrast and cross-modal contrast. They are formulated as follows.

The Intra-Modal Contrastive (IMC) Objective injects ViPFormer with the ability to resist data transformations and small perturbations (e.g., translation, rotation, jittering) to the same objects while maximizing the distance of different objects in feature space. This strategy will make the pre-trained model insensitive to random noises and generalize better. Specifically, a point cloud P is transformed by two data augmentations t_1 and t_2 , resulting in P^{t_1} and P^{t_2} . After going through ViPFormer, their feature representations $\mathbf{p}^{t_1} = \mathbf{o}_P^{t_1}$ and $\mathbf{p}^{t_2} = \mathbf{o}_P^{t_2}$ are obtained. The IMC objective \mathcal{L}_{imc} is formulated by NT-Xent loss [54]:

$$l(i, t_1, t_2) = -\log \frac{\exp(s(\mathbf{p}_i^{t_1}, \mathbf{p}_i^{t_2})/\tau)}{\sum_{\substack{k=1 \\ k \neq i}}^N \exp(s(\mathbf{p}_i^{t_1}, \mathbf{p}_k^{t_1})/\tau) + \sum_{k=1}^N \exp(s(\mathbf{p}_i^{t_1}, \mathbf{p}_k^{t_2})/\tau)} \quad (8)$$

$$\mathcal{L}_{imc} = \frac{1}{2N} \sum_{i=1}^N (l(i, t_1, t_2) + l(i, t_2, t_1)) \quad (9)$$

where N is the batch size, τ is the temperature coefficient and $s(\cdot)$ represents the cosine similarity.

The Cross-Modal Contrastive (CMC) Objective maximizes the agreement of feature representations of paired images and point clouds, while minimizing that of unpaired ones in the same feature space. ViPFormer achieves this goal only when it understands the information contained in both

modalities. Similarly, the CMC objective \mathcal{L}_{cmc} is formulated by NT-Xent loss [54]:

$$l(i, \mathbf{p}, \mathbf{f}) = -\log \frac{\exp(s(\mathbf{p}_i, \mathbf{f}_i)/\tau)}{\sum_{\substack{k=1 \\ k \neq i}}^N \exp(s(\mathbf{p}_i, \mathbf{p}_k)/\tau) + \sum_{k=1}^N \exp(s(\mathbf{p}_i, \mathbf{f}_k)/\tau)} \quad (10)$$

$$\mathcal{L}_{cmc} = \frac{1}{2N} \sum_{i=1}^N (l(i, \mathbf{p}, \mathbf{f}) + l(i, \mathbf{f}, \mathbf{p})) \quad (11)$$

N , τ and $s(\cdot)$ have the same meaning as those in Eq. 8.

During pretraining, ViPFormer combines IMC and CMC as the overall loss. A balancing factor α is deployed between two objectives as the cross-modal loss \mathcal{L}_{cmc} is harder to optimize and usually several times bigger than the intra-modal loss \mathcal{L}_{imc} .

$$\mathcal{L} = \mathcal{L}_{imc} + \alpha \mathcal{L}_{cmc} \quad (12)$$

IV. EXPERIMENTS

In this section, firstly, we elaborate on the pretraining settings of ViPFormer. Then the pretrained ViPFormer is transferred to various downstream tasks to evaluate its performances. Thirdly, the effectiveness of different components in ViPFormer is validated by extensive ablation studies. Finally, the predictions of ViPFormer on different tasks are visualized for a better understanding.

A. Pretraining Settings

Datasets We use the same dataset as in [22]. The point clouds and images come from ShapeNet [55] and DISN [56]. There are 43,783 point clouds and each point cloud P corresponds to 24 rendered images, from which an image I is randomly selected to pair with P . During pretraining, point cloud P contains 2048 points and the corresponding image I is resized to $144 \times 144 \times 3$. In FPS and k NN, the point cloud is divided into $G=128$ centers and $k=32$ nearest neighbors are retrieved for each center. The image patch size Q is set to 12.

Architecture In Input Adapter, the dimension of point/image patch embedding is projected to 384. In Encoder, there are $L=9$ stacked MSA and MLP layers. All MSA layers have 6 heads. The widening ratio of the MLP hidden layer is 4. In Output Adapter, the 2-layer MLP is of size $\{768, 384, 384\}$. We justify the design choices of the architecture through controlled experiments in Section IV-C.1.

Optimization We pretrain ViPFormer for 300 epochs, adopting AdamW [57] as the optimizer and CosineAnnealingWarmupRestarts [58] as the learning rate scheduler. The restart interval is 100 epochs and the warmup period is the first 5 epochs. The learning rate scales linearly to peak during each warmup, then decays with the cosine annealing schedule. The initial learning rate peak is 0.001, multiplied by 0.6 after each interval. The balancing factor α is set to 1, which works well. We record the best pretrained model according to the zero-shot accuracy on ModelNet40 [59].

B. Model Complexity, Latency and Performance on Downstream Tasks

In this part, the pretrained ViPFormer is transferred to various downstream tasks to evaluate its complexity, latency and performance. These metrics are critical dimensions for assessing point cloud understanding methods. Complexity is reflected by a model’s number of parameters (#Params). Latency is counted by running time and performance is subject to overall accuracy (OA) in the classification task and mean class-wise Intersection of Union (mIoU) in the segmentation task. We compare with previous state-of-the-art unsupervised methods.

Point Cloud Classification The experiments are conducted on two widely used datasets: ScanObjectNN [60] and ModelNet40 [59]. ScanObjectNN contains 2880 objects from 15 categories. It is challenging because objects in this dataset are usually cluttered with background or are partial due to occlusions. ModelNet40 is a synthetic point cloud dataset, including 12308 objects across 40 categories. We use the same settings as previous work [19], [20], [14], [22] to sample 1024 points to represent a 3D object. We reimplement previous methods according to the released codes since they do not report the #Params and latency. For latency, we consider two stages (Pretrain and Finetune) and count the running time of a single epoch in each stage.

For the ScanObjectNN [60] dataset, all methods are finetuned on it and the results are recorded in Tab. I. The best score is in bold black and the second best score is marked in blue. ViPFormer not only outperforms previous state-of-the-art Point-MAE by 0.7% in classification accuracy but also reduces 76.9% #Params and runs $\sim 2.6x$ faster than Point-MAE.

TABLE I: Comparison of model complexity, latency and performance with existing unsupervised methods on ScanObjectNN.

Method	#Params	Latency		OA
	(M)	Pretrain (s)	Finetune (ms)	(%)
FoldingNet [10]	2.0	–	–	81.0
PointContrast [19]	37.9	–	–	79.6
DepthContrast [20]	8.2	–	–	80.4
OcCo [14]	3.5	~ 600.0	16,100	83.3
CrossPoint [22]	27.7	946.0	14,000	81.7
Point-BERT [11]	39.1	633.5	3,973	87.4
Point-MAE [12]	22.1	576.0	3,612	90.0
ViPFormer	5.1	22.2	1,015	90.7

The classification results on ModelNet40 are shown in Table II. The Pretrain latency is not changed because pretraining is independent of the downstream datasets, including ScanObjectNN and ModelNet40. ModelNet40 is a larger dataset so finetuning on it consumes more time. ViPFormer achieves higher classification accuracy with lower model complexity and runtime latency. It leads Point-MAE by 0.7% accuracy while reducing #Params by 24.1%.

Object Part Segmentation We also transfer ViPFormer to the task of object part segmentation. The experiments

TABLE II: Comparison of model complexity, latency and performance with existing unsupervised methods on ModelNet40.

Method	#Params (M)	Latency		OA (%)
		Pretrain (s)	Finetune (ms)	
FoldingNet [10]	2.0	–	–	90.6
PointContrast [19]	37.9	–	–	90.0
DepthContrast [20]	8.2	–	–	89.2
OcCo [14]	3.5	~600.0	39,295	92.5
CrossPoint [22]	27.7	946.0	35,258	90.3
Point-BERT [11]	39.1	633.5	10,329	93.0
Point-MAE [12]	22.1	576.0	9,344	93.2
ViPFormer	16.7	60.9	4,198	93.9

are conducted on the ShapeNetPart [61] dataset which contains 16881 point clouds and each point cloud consists of 2048 points. Objects in ShapeNetPart are divided into 16 categories and 50 annotated parts. For a fair comparison, we follow previous work [11] [12] to add a simple part segmentation head on ViPFormer Encoder. The pretrained weights of ViPFormer are used to initialize the part segmentation model. In addition to the above metrics, mean class-wise IoU (mIoU) is added to evaluate the part segmentation performance. The results are reported in Tab. III. ViPFormer reaches comparable OA and mIoU with best performing Point-MAE while having lower model complexity.

TABLE III: Object part segmentation on ShapeNetPart.

Method	#Params (M)	Latency (s)	OA (%)	mIoU (%)
PointContrast [19]	37.9	–	–	–
OcCo [14]	1.5	32.0	93.9	79.7
CrossPoint [22]	27.5	80.0	93.8	84.3
Point-BERT [11]	44.1	58.8	–	84.1
Point-MAE [12]	27.1	46.3	94.8	84.7
ViPFormer	26.8	42.1	94.8	84.7

Few-shot Object Classification Few-shot evaluation is used to validate the transferability of a pretrained model with limited labeled data. The conventional setting is “ N -way, K -shot” [20], [14], [22]. Under this setting, N classes and K samples in a downstream task dataset are randomly selected for training an SVM model of the linear kernel. The test score on the downstream task given by SVM can reflect the quality of the pretrained model as the inputs to the SVM model are the features extracted by the pretrained model. Here the downstream task datasets are ModelNet40 and ScanObjectNN, respectively. We perform 10 runs for each “ N -way, K -shot” and compute their mean and standard deviation. The results are shown in Table IV. On ModelNet40, ViPFormer achieves comparable accuracy with previous strong baselines, whereas it shows consistent improvements on ScanObjectNN. The IMC and CMC objectives enable ViPFormer to understand the information contained in both modalities, so it can better deal with the challenging scenarios in ScanObjectNN.

TABLE IV: Comparison of few-shot classification accuracy with existing methods on ModelNet40 and ScanObjectNN.

Method	5-way		10-way	
	10-shot	20-shot	10-shot	20-shot
ModelNet40				
OcCo [14]	90.6±2.8	92.5±1.9	82.9±1.3	86.5±2.2
CrossPoint [22]	91.0±2.9	95.0±3.4	82.2±6.5	87.8±3.0
ViPFormer	91.1±7.2	93.4±4.5	80.8±4.2	87.1±5.8
ScanObjectNN				
OcCo [14]	72.4±1.4	77.2±1.4	57.0±1.3	61.6±1.2
CrossPoint [22]	72.5±8.3	79.0±1.2	59.4±4.0	67.8±4.4
ViPFormer	74.2±7.0	82.2±4.9	63.5±3.8	70.9±3.7

C. Ablation Studies

Ablation studies are conducted to 1) justify the architecture of ViPFormer, 2) demonstrate the effectiveness of IMC and CMC optimization objectives, and 3) analyze the advantages of the pretrain-finetune strategy over training from scratch.

1) *Architecture*: The controlled variables of ViPFormer architecture are the number of self-attention layers (#SA_Layers), the widening ratio of the MLP hidden layer (MLP ratio), the number of attention heads (#Heads), the sequence length (#Length) and the model dimension (D_model). For different architectures, the accuracy of the pretrain-finetune scheme is reported on ModelNet40 and ScanObjectNN, respectively, shown in Tab. V. The overall trend is the larger the model, the better the performance. We choose the best-performing architecture to compare with other methods.

2) *Contrastive Optimization Objectives*: The effectiveness of proposed IMC and CMC contrastive objectives are evaluated by training ViPFormer in three modes: i) only use IMC, ii) only use CMC, and iii) use IMC and CMC together. The experiments are conducted on different learning stages (Pretrain vs. Finetune) and different datasets (ModelNet40 vs. ScanObjectNN). The results are shown in Tab. VI. Apparently, the combination of IMC and CMC optimization objectives significantly improves the model performance for target tasks across different datasets.

3) *Learning Strategies*: The differences between two kinds of learning strategies, Train from scratch and Pretrain-Finetune, are also investigated. As Tab. VII shows, The Pretrain-Finetune strategy outperforms Train from scratch by 1.9% and 4.1% on ModelNet40 and ScanObjectNN, respectively. The results indicate the initialization provided by the pretrained ViPFormer really helps the model find better directions and solutions in downstream tasks.

D. Visualization

Object Part Segmentation We conduct experiments on ShapeNetPart [61] to visualize the predictions of ViPFormer to different object parts. This dataset has 16 object categories and we randomly select one object from each category. ViPFormer predicts part labels for all points in the selected object. Then different part labels are mapped to different colors. As Fig. 3 presents, ViPFormer successfully handles different objects and segments their parts in most cases.

TABLE V: Ablation Study: Model Architecture.

#SA_Layers	7	7	9	9	7	7	9	9	7	7	9	9	7	7	9	9
MLP ratio	2	2	2	2	4	4	4	4	2	2	2	2	4	4	4	4
#Heads	4	4	4	4	4	4	4	4	6	6	6	6	6	6	6	6
#Length	96	128	96	128	96	128	96	128	96	128	96	128	96	128	96	128
D_model	256	256	256	256	256	256	256	256	384	384	384	384	384	384	384	384
Accuracy																
ModelNet40	91.5	91.5	91.5	92.5	93.2	91.0	92.7	91.3	93.0	93.2	92.0	92.2	92.2	92.2	93.2	93.9
ScanObjectNN	84.5	83.5	86.6	90.7	85.6	85.6	86.6	87.6	88.7	84.5	84.5	89.7	89.7	88.7	89.7	89.7

TABLE VI: Ablation Study: Performance comparison on ModelNet40 (M) and ScanObjectNN (S) when using different contrastive objectives.

Modality Types	Pretrain		Finetune	
	OA _M	OA _S	OA _M	OA _S
IMC <i>only</i>	86.4	76.4	91.3	87.6
CMC <i>only</i>	87.3	66.4	91.3	81.4
IMC & CMC	87.0	75.7	93.9	89.7

TABLE VII: Ablation Study: Comparison of the classification accuracy using Pretrain-Finetune and Train-from-scratch strategy. The used datasets are ModelNet40 (M) and ScanObjectNN (S).

Learning Strategy	OA _M	OA _S
Train from scratch	92.0	85.6
Pretrain-Finetune	93.9	89.7




Fig. 3: Object part segmentation predictions of ViPFormer

Feature Distribution The distributions of pretrained and finetuned features are visualized by t-SNE [62], exhibited in Fig. 4. The experiments are conducted on ModelNet40 and ScanObjectNN. The pretrained features roughly scatter into different locations and provide good initialization for downstream tasks. After finetuning on the target datasets, the features are clearly separated by different clusters.

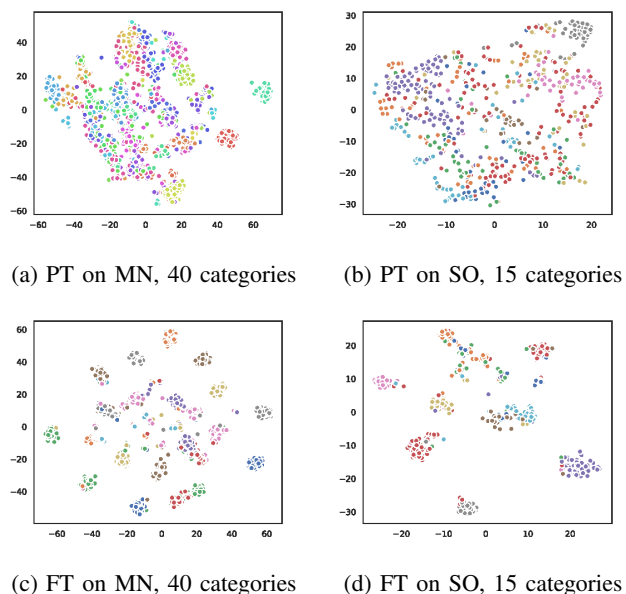


Fig. 4: t-SNE [62] Visualization of pretrained (PT) and finetuned (FT) features on ModelNet40 (MN) and ScanObjectNN (SO).

V. CONCLUSION

In this paper, We propose an efficient Vision-and-Pointcloud Transformer to unify image and point cloud processing in a single architecture. ViPFormer is pretrained by optimizing intra-modal and cross-modal contrastive objectives. When transferred to downstream tasks and compared with existing unsupervised methods, ViPFormer shows advantages in model complexity, runtime latency and performances. And the contribution of each component is validated by extensive ablation studies. In the future, we should pay more attention to the image branch and explore its performances on downstream tasks since the current version focuses on point cloud understanding.

REFERENCES

- [1] T. Xiang, C. Zhang, Y. Song, J. Yu, and W. Cai, "Walk in the cloud: Learning curves for point clouds shape analysis," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 915–924.
- [2] X. Ma, C. Qin, H. You, H. Ran, and Y. Fu, "Rethinking network design and local geometry in point cloud: A simple residual MLP framework," in *International Conference on Learning Representations*, 2022. [Online]. Available: https://openreview.net/forum?id=3Pbra-_u76D
- [3] Q. Hu, B. Yang, L. Xie, S. Rosa, Y. Guo, Z. Wang, N. Trigoni, and A. Markham, "Randla-net: Efficient semantic segmentation of large-scale point clouds," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [4] H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun, "Point transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 16 259–16 268.
- [5] C. Park, Y. Jeong, M. Cho, and J. Park, "Fast point transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 16 949–16 958.
- [6] B. Yang, W. Luo, and R. Urtaşun, "Pixor: Real-time 3d object detection from point clouds," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [7] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [8] C. R. Qi, O. Litany, K. He, and L. J. Guibas, "Deep hough voting for 3d object detection in point clouds," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [9] Y. Wang, T. Ye, L. Cao, W. Huang, F. Sun, F. He, and D. Tao, "Bridged transformer for vision and point cloud 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 12 114–12 123.
- [10] Y. Yang, C. Feng, Y. Shen, and D. Tian, "Foldingnet: Point cloud auto-encoder via deep grid deformation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [11] X. Yu, L. Tang, Y. Rao, T. Huang, J. Zhou, and J. Lu, "Point-bert: Pre-training 3d point cloud transformers with masked point modeling," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 19 313–19 322.
- [12] Y. Pang, W. Wang, F. E. H. Tay, W. Liu, Y. Tian, and L. Yuan, "Masked autoencoders for point cloud self-supervised learning," in *Computer Vision – ECCV 2022*. Springer International Publishing, 2022.
- [13] J. Sauder and B. Sievers, "Self-supervised deep learning on point clouds by reconstructing space," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/993edc98ca87f7e08494ecc37fa836f7-Paper.pdf>
- [14] H. Wang, Q. Liu, X. Yue, J. Lasenby, and M. J. Kusner, "Unsupervised point cloud pre-training via occlusion completion," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 9782–9792.
- [15] X. Yu, Y. Rao, Z. Wang, Z. Liu, J. Lu, and J. Zhou, "PointR: Diverse point cloud completion with geometry-aware transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 12 498–12 507.
- [16] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum, "Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016. [Online]. Available: <https://proceedings.neurips.cc/paper/2016/file/44f683a84163b3523afe57c2e008bc8c-Paper.pdf>
- [17] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas, "Learning representations and generative models for 3D point clouds," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 40–49. [Online]. Available: <https://proceedings.mlr.press/v80/achlioptas18a.html>
- [18] Z. Han, M. Shang, Y.-S. Liu, and M. Zwicker, "View inter-prediction gan: Unsupervised representation learning for 3d shapes by learning global shape memories to support local view predictions," in *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, ser. AAAI'19/IAAI'19/EAAI'19. AAAI Press, 2019. [Online]. Available: <https://doi.org/10.1609/aaai.v33i01.33018376>
- [19] S. Xie, J. Gu, D. Guo, C. R. Qi, L. Guibas, and O. Litany, "Pointcontrast: Unsupervised pre-training for 3d point cloud understanding," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 574–591.
- [20] Z. Zhang, R. Girdhar, A. Joulin, and I. Misra, "Self-supervised pretraining of 3d features on any point-cloud," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 10 252–10 263.
- [21] J. Hou, B. Graham, M. Nießner, and S. Xie, "Exploring data-efficient 3d scene understanding with contrastive scene contexts," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 587–15 597.
- [22] M. Afham, I. Dissanayake, D. Dissanayake, A. Dharmasiri, K. Thilakarathna, and R. Rodrigo, "Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 9902–9912.
- [23] L. Tang, Y. Zhan, Z. Chen, B. Yu, and D. Tao, "Contrastive boundary learning for point cloud segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 8489–8499.
- [24] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 16 000–16 009.
- [25] W. Kim, B. Son, and I. Kim, "Vilt: Vision-and-language transformer without convolution or region supervision," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 5583–5594. [Online]. Available: <http://proceedings.mlr.press/v139/kim21k.html>
- [26] K. Desai and J. Johnson, "VirTex: Learning Visual Representations from Textual Annotations," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Nashville, TN, USA: IEEE, June 2021, pp. 11 157–11 168. [Online]. Available: <https://ieeexplore.ieee.org/document/9577368/>
- [27] M. B. Sariyildiz, J. Perez, and D. Larlus, "Learning Visual Representations with Caption Annotations," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, vol. 12353, pp. 153–170, series Title: Lecture Notes in Computer Science. [Online]. Available: https://link.springer.com/10.1007/978-3-030-58598-3_10
- [28] A. Owens and A. A. Efros, "Audio-Visual Scene Analysis with Self-Supervised Multisensory Features," in *Computer Vision – ECCV 2018*, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds. Cham: Springer International Publishing, 2018, vol. 11210, pp. 639–658, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-030-01231-1_39
- [29] P. Morgado, N. Vasconcelos, and I. Misra, "Audio-Visual Instance Discrimination with Cross-Modal Agreement," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Nashville, TN, USA: IEEE, June 2021, pp. 12 470–12 481. [Online]. Available: <https://ieeexplore.ieee.org/document/9578129/>
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

- [31] C. R. Qi, H. Su, K. Mo, and L. 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 (CVPR)*, July 2017.
- [32] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *ACM Transactions on Graphics (TOG)*, 2019.
- [33] P. Jiang and S. Saripalli, "Contrastive Learning of Features between Images and LiDAR," arXiv, Tech. Rep. arXiv:2206.12071, June 2022, arXiv:2206.12071 [cs] type: article. [Online]. Available: <http://arxiv.org/abs/2206.12071>
- [34] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.
- [35] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/d8bf84be3800d12f74d8b05e9b89836f-Paper.pdf>
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [37] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=YicbFdNTTy>
- [38] H. Bao, L. Dong, S. Piao, and F. Wei, "BEiT: BERT pre-training of image transformers," in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=p-BhZSz59o4>
- [39] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis: Association for Computational Linguistics, June 2019, pp. 4171–4186. [Online]. Available: <https://aclanthology.org/N19-1423>
- [40] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," 2019.
- [41] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 1877–1901. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>
- [42] X. Pan, Z. Xia, S. Song, L. E. Li, and G. Huang, "3d object detection with pointformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 7463–7472.
- [43] C. Sharma and M. Kaul, "Self-supervised few-shot learning on point clouds," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 7212–7221. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/50c1f44e426560f3f2cdbc3e19e39903-Paper.pdf>
- [44] L. Nunes, R. Marcuzzi, X. Chen, J. Behley, and C. Stachniss, "SegContrast: 3D Point Cloud Feature Representation Learning through Self-supervised Segment Discrimination," *ral*, vol. 7, no. 2, pp. 2116–2123, 2022. [Online]. Available: <http://www.ipb.uni-bonn.de/pdfs/nunes2022ral-icra.pdf>
- [45] X. Ma, C. Qin, H. You, H. Ran, and Y. Fu, "Rethinking network design and local geometry in point cloud: A simple residual MLP framework," in *International Conference on Learning Representations*, 2022. [Online]. Available: https://openreview.net/forum?id=3Pbra_u76D
- [46] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I*. Berlin, Heidelberg: Springer-Verlag, 2020, p. 213–229. [Online]. Available: https://doi.org/10.1007/978-3-030-58452-8_13
- [47] I. Misra, R. Girdhar, and A. Joulin, "An end-to-end transformer model for 3d object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 2906–2917.
- [48] J. Mao, Y. Xue, M. Niu, H. Bai, J. Feng, X. Liang, H. Xu, and C. Xu, "Voxel transformer for 3d object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 3164–3173.
- [49] H. Sheng, S. Cai, Y. Liu, B. Deng, J. Huang, X.-S. Hua, and M.-J. Zhao, "Improving 3d object detection with channel-wise transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 2743–2752.
- [50] Z. Liu, Z. Zhang, Y. Cao, H. Hu, and X. Tong, "Group-free 3d object detection via transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 2949–2958.
- [51] M.-H. Guo, J.-X. Cai, Z.-N. Liu, T.-J. Mu, R. R. Martin, and S.-M. Hu, "Pct: Point cloud transformer," 2020.
- [52] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira, "Perceiver: General perception with iterative attention," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 4651–4664. [Online]. Available: <https://proceedings.mlr.press/v139/jaegle21a.html>
- [53] A. Jaegle, S. Borgeaud, J.-B. Alayrac, C. Doersch, C. Ionescu, D. Ding, S. Koppula, D. Zoran, A. Brock, E. Shelhamer, O. J. Henaff, M. Botvinick, A. Zisserman, O. Vinyals, and J. Carreira, "Perceiver IO: A general architecture for structured inputs & outputs," in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=fILj7WpI-g>
- [54] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proceedings of the 37th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 1597–1607. [Online]. Available: <https://proceedings.mlr.press/v119/chen20j.html>
- [55] A. X. Chang, T. A. Funkhouser, L. J. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "Shapenet: An information-rich 3d model repository," *CoRR*, vol. abs/1512.03012, 2015. [Online]. Available: <http://arxiv.org/abs/1512.03012>
- [56] Q. Xu, W. Wang, D. Ceylan, R. Mech, and U. Neumann, "Disn: Deep implicit surface network for high-quality single-view 3d reconstruction," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/39059724f73a9969845dfe4146c5660e-Paper.pdf>
- [57] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=Bkg6RiCqY7>

- [58] N. Katsura, "Pytorch cosineannealing with warmup restarts," <https://github.com/katsura-jp/pytorch-cosine-annealing-with-warmup>, 2021.
- [59] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3d shapenets: A deep representation for volumetric shapes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [60] M. A. Uy, Q.-H. Pham, B.-S. Hua, T. Nguyen, and S.-K. Yeung, "Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [61] L. Yi, V. G. Kim, D. Ceylan, I.-C. Shen, M. Yan, H. Su, C. Lu, Q. Huang, A. Sheffer, and L. Guibas, "A scalable active framework for region annotation in 3d shape collections," *ACM Trans. Graph.*, vol. 35, no. 6, nov 2016. [Online]. Available: <https://doi.org/10.1145/2980179.2980238>
- [62] L. van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of Machine Learning Research*, vol. 9, no. 86, pp. 2579–2605, 2008. [Online]. Available: <http://jmlr.org/papers/v9/vandermaaten08a.html>