

DiffCLIP: Leveraging Stable Diffusion for Language Grounded 3D Classification

Sitian Shen¹, Zilin Zhu¹, Linqian Fan², Harry Zhang³, Xinxiao Wu¹

¹Beijing Key Laboratory of Intelligent Information Technology, School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China,

²Beijing Jiaotong University, Beijing, China,

³Massachusetts Institute of Technology, Cambridge, USA,

sitian@bit.edu.cn, smzzlcn@gmail.com, LinqFan@bjtu.edu.cn, harryz@mit.edu, wuxinxiao@bit.edu.cn

Abstract

Large pre-trained models have revolutionized the field of computer vision by facilitating multi-modal learning. Notably, the CLIP model has exhibited remarkable proficiency in tasks such as image classification, object detection, and semantic segmentation. Nevertheless, its efficacy in processing 3D point clouds is restricted by the domain gap between the depth maps derived from 3D projection and the training images of CLIP.

This paper introduces DiffCLIP, a novel pre-training framework that seamlessly integrates stable diffusion with ControlNet. The primary objective of DiffCLIP is to bridge the domain gap inherent in the visual branch. Furthermore, to address few-shot tasks in the textual branch, we incorporate a style-prompt generation module.

Extensive experiments on the ModelNet10, ModelNet40, and ScanObjectNN datasets show that DiffCLIP has strong abilities for 3D understanding. By using stable diffusion and style-prompt generation, DiffCLIP achieves an accuracy of 43.2% for zero-shot classification on OBJ_BG of ScanObjectNN, which is state-of-the-art performance, and an accuracy of 82.4% for zero-shot classification on ModelNet10, which is also state-of-the-art performance.

1. Introduction

Significant advancements in Natural Language Processing (NLP), particularly in the context of large pre-trained models, have considerably impacted computer vision, enabling multi-modal learning, transfer learning, and spurring the development of new architectures [6]. These pre-trained NLP models have been instrumental in the extraction of high-level features from text data, subsequently enabling the fusion with visual features from images or videos, which in turn improves performance on tasks such as image captioning and video classification. Recently, OpenAI's Con-

trastive Language Image Pre-training (CLIP) [19] model has been particularly popular due to its superior performance in image classification, object detection, and semantic segmentation. These achievements underscore the potential of large vision-language pre-trained models across a spectrum of applications.

The remarkable performance of large pre-trained models on 2D vision tasks has stimulated researchers to investigate the potential applications of these models within the domain of 3D point cloud processing. Several approaches, including PointCLIP [39], PointCLIP v2 [43] and CLIP2Point [11] have been proposed over the years. While these methods demonstrate advantages in zero-shot and few-shot 3D point cloud processing tasks, experimental results on real-world datasets suggest that their efficacy in handling real-world tasks is limited. This limitation is primarily attributed to the significant domain gap between 2D depth maps derived from 3D projection and training images of CLIP. Consequently, the primary focus of our work is to minimize this domain gap and improve the performance of CLIP on zero-shot and few-shot 3D vision tasks.

Previous studies have made significant progress in addressing this challenge. For instance, [10] leverages prompt learning to adapt the input domain of the pre-trained visual encoder from computer-generated images of CAD objects to real-world images. The CLIP² [36] model aligns three-dimensional space with a natural language representation that is applicable in real-world scenarios, facilitating knowledge transfer between domains without prior training. To further boost the performance, we introduce a style feature extraction step and a style transfer module via stable diffusion [22] to further mitigate the domain gap.

In this paper, we propose a novel pre-training framework, called DiffCLIP, to minimize the domain gap within both the visual and textual branches. On the visual side, we design a multi-view projection on the original 3D point cloud, resulting in multi-view depth maps. Each depth map undergoes a style transfer process guided by stable diffusion

and ControlNet [37], generating photorealistic 2D RGB images. The images generated from the style transfer module are then input into the frozen image encoder of CLIP. For zero-shot tasks, we design handcrafted prompts for both the stable diffusion module and the frozen text encoder of CLIP. For few-shot tasks, we further establish a style-prompt generation module that inputs style features extracted from a pre-trained point-cloud-based network and processed via a meta-net, and outputs the projected features as the style-prompt. Coupled with the class label, the entire prompt is then fed into the frozen text encoder of CLIP to guide the downstream process.

To demonstrate the effectiveness of DiffCLIP, we evaluate its zero-shot and few-shot classification ability on both synthetic and real-world datasets like ModelNet and ScanObjectNN. Our main contributions are summarized as follows:

1. We propose DiffCLIP, a novel neural network architecture that combines a pre-trained CLIP model with point-cloud-based networks for enhanced 3D understanding.
2. We develop a technique that effectively minimizes the domain gap in point-cloud processing, leveraging the capabilities of stable diffusion and style-prompt generation, which provides significant improvements in the task of point cloud understanding.
3. We conduct experiments on widely adopted benchmark datasets such as ModelNet and the more challenging ScanObjectNN to illustrate the robust 3D understanding capability of DiffCLIP and conduct ablation studies that evaluate the significance of each component within DiffCLIP’s architecture.

2. Related Work

2.1. 3D Shape Classification

3D point cloud processing methods for classification can be divided into three categories: projection-based [13; 33], volumetric-based [14; 21] and point-based [8; 40]. In projection-based methods, 3D shapes are often projected into multiple views to obtain comprehensive features. Then they are fused with well-designed weight. The most common instances of such a method are MVCNN [27] which employs a straightforward technique of max-pooling to obtain a global descriptor from multi-view features, and MHBN [34] that makes use of harmonized bilinear pooling to integrate local convolutional features and create a smaller-sized global descriptor. Among all the point-based methods, PointNet [16] is a seminal work that directly takes point clouds as its input and achieves permutation invariance with a symmetric function. Since features are learned independently for each point in PointNet [16], the local structural information between points is hardly captured. Therefore, PointNet++ [17]

is proposed to capture fine geometric structures from the neighborhood of each point.

Recently, several 3D large pre-trained models are proposed. CLIP [19] is a cross-modal text-image pre-training model based on contrastive learning. PointCLIP [39] and PointCLIP v2 [43] extend CLIP’s 2D pre-trained knowledge to 3D point cloud understanding. Some other 3D processing methods based on large pre-trained model CLIP are proposed in recent years, including CLIP2Point [11], CLIP² [36], and CG3D [10]. Similar to some of those methods, our model combines widely used point cloud learning architectures with novel large pre-trained models such as CLIP. The key difference here is that our point-based network hierarchically guides the generation of prompts by using direct 3D point cloud features, which largely reduces the domain gap between training and testing data for better zero-shot and few-shot learning. We also design a simpler yet more effective projection method, an intuitive multi-view fusion strategy, to make the projection processing as efficient as possible for the zero-shot task.

2.2. Domain Adaptation and Generalization in Vision-Language Models

Vision-language pre-trained models such as CLIP exhibit impressive zero-shot generalization ability to the open world. Despite the good zero-shot performance, it is found that further adapting CLIP using task-specific data comes at the expense of out-of-domain (OOD) generalization ability. [20; 31]. Recent advances explore improving the OOD generalization of CLIP on the downstream tasks by adapter learning [7; 38], model ensembling [31], test-time adaptation [25], prompt learning [41; 42], and model fine-tuning [26]. Specifically, in StyLIP [1], style features are extracted hierarchically from the image encoder of CLIP to generate domain-specific prompts. To the best of our knowledge, our work is the first to extract style features of 3D point clouds directly from point-based processing networks. With the extracted style features, we are able to generate domain-specific prompts for the text encoder of CLIP.

2.3. Style Transfer using Diffusion Model

Traditional style transfer is one of the domain generalization methods which helps reduce the domain gap between source and target domains by aligning their styles while preserving the content information [12]. For example, Domain-Specific Mappings [4] has been proposed to transfer the style of source domains to a target domain without losing domain-specific information. Fundamentally different from traditional methods of style transfer that operate on pre-existing images, text-to-image diffusion models generate an image based on a textual description that may include style-related information, thus embedding the style defined in textual descriptions into the generated images. These models are based

on a diffusion process, which involves gradually smoothing out an input image to create a stylized version of it, showing a strong ability in capturing complex texture and style information in images. For instance, Glide [15], Disco Diffusion [5], Stable Diffusion [23], and Imagen [24] all support text-to-image generation and editing via a diffusion process, which can further be used to change the image styles via textual prompts variations. We refer to the textual prompts in diffusion models as **style prompts** as they often encapsulate the style of generated images. In our work, we implement different text-to-image diffusion models and choose ControlNet [37] to help transfer depth map from 3D projection to a more realistic image style, minimizing its style gap with the training images of CLIP.

3. Method

DiffCLIP is a novel neural network architecture that enhances the performance of the CLIP model on 3D point processing tasks via style transfer and style-prompt generation. In Section 3.1, we first revisit PointCLIP [39] and PointCLIP v2 [43] that first introduce the CLIP model into the processing of 3D point clouds. Then in Section 3.2, we introduce the framework and motivation of each part of DiffCLIP, which minimizes the domain gap between 3D tasks and 2D pre-training data. In Section 3.3, we present the detailed usage of DiffCLIP on zero-shot and few-shot classification tasks.

3.1. Revisiting PointCLIP and PointCLIP v2

PointCLIP [39] is a recent extension of CLIP, which enables zero-shot classification on point cloud data. It involves taking a point cloud and converting it into multiple depth maps from different angles without rendering. The predictions from the individual views are combined to transfer knowledge from 2D to 3D. Moreover, an inter-view adapter is used to improve the extraction of global features and fuse the few-shot knowledge acquired from 3D into CLIP that was pre-trained on 2D image data.

Specifically, for an unseen dataset of K classes, PointCLIP constructs the textual prompt by placing all category names into a manual template, termed prompt, and then encodes the prompts into a D -dimensional textual feature, acting as the zero-shot classifier $W_t \in \mathbb{R}^{K \times D}$. Meanwhile, the features of each projected image from M views are encoded as f_i for $i = 1, \dots, M$ by the visual encoder of CLIP. On top of this, the classification logits s_i of view i and the final logits s_p of each point are calculated by

$$\text{logits}_i = f_i W_t^T, \forall i = 1, \dots, M \quad (1a)$$

$$\text{logit}_{s_p} = \sum_{i=1}^M \alpha_i \text{logits}_i, \quad (1b)$$

where α_i is a hyper-parameter weighing the importance of view i .

PoinCLIP v2 [43] introduces a realistic shape projection module to generate more realistic depth maps for the visual encoder of CLIP, further narrowing down the domain gap between projected point clouds and natural images. Moreover, a large-scale language model, GPT-3 [2], is leveraged to automatically design a more descriptive 3D-semantic prompt for the textual encoder of CLIP instead of the previous hand-crafted one.

3.2. DiffCLIP Framework

The DiffCLIP framework is illustrated in Fig. 1. We describe each component of DiffCLIP in detail in the following sections.

3.2.1 Multi-View Realistic Projection

In DiffCLIP, we use stable diffusion [22] to help transfer projected depth maps to a more realistic image style, minimizing the domain gap with the training images of CLIP. To generate realistic depth maps for better controlling the stable diffusion model, as well as to save computational resources for zero-shot and few-shot tasks, we design a multi-view realistic projection module, which has three steps: proportional sampling, central projection, and 2D max-pooling densifying.

Portional Sampling. We randomly sample points on each face and edge of the object in the dataset, the number of which is proportional to the area of the faces and the length of the edges. Specifically, assuming that k_l points are sampled on a line with length l and k_f points are sampled on a face with area s , we set sampling threshold values β_1 and β_2 , so the number of sampled points can be calculated by $k_l = \frac{l}{\beta_1}$ and $k_f = \frac{s}{\beta_2}$.

Central Projection. For each projection viewpoint $n \in \{1, 2, \dots, N\}$, we select a specific projection center and projection plane, and use affine transformation to calculate the coordinates of the projection point on the projection plane for each sampling point. The pixel intensity d is calculated as $d = \frac{1}{dis}$, where dis is the distance between the sampling point and our projection plane. On the 2D plane, the projected point may not fall exactly on a pixel, so we use the nearest-neighbor pixel to approximate its density.

2D Max-pooling Densifying. In the central projection step, the nearest-neighbor pixel approximation may yield many pixels to be unassigned, causing the projected depth map to look unrealistic, sparse, and scattered. To tackle this problem, we densify the projected points via a local max-pooling operation to guarantee visual continuity. For each projected pixel of the depth map, we choose the max density d_{max} among four points: the pixel itself and the pixels to its right, bottom, and bottom right. We then assign the max density to the neighboring four pixels.

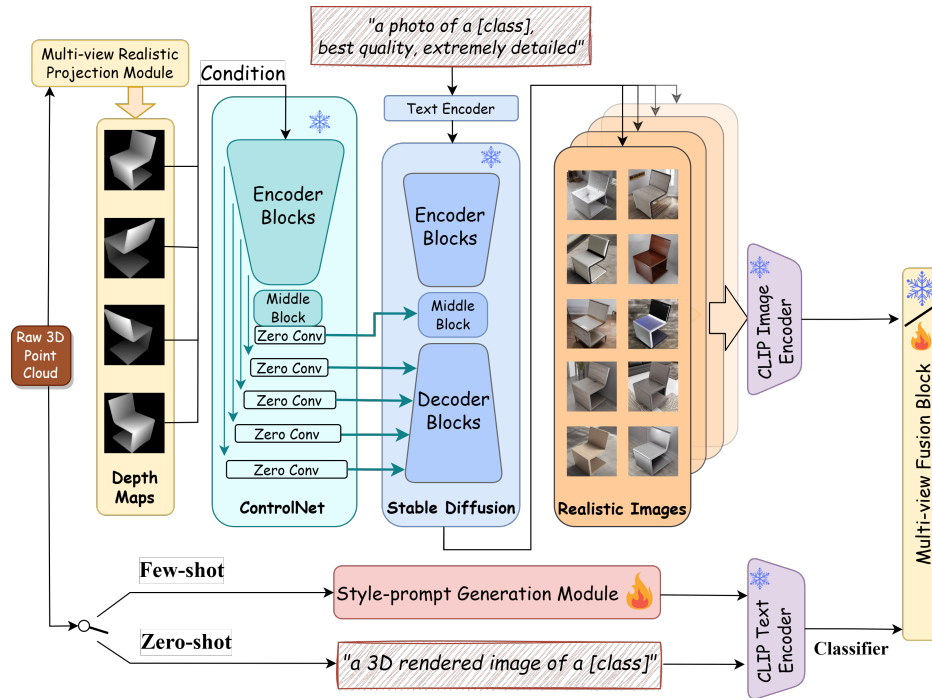


Figure 1. **Framework Structure of DiffCLIP.** In the visual branch, DiffCLIP has two modules: Multi-View Realistic Projection Module, which produces multi-view depth maps, and Stable-Diffusion-Based Style Transfer Module, which uses a pre-trained ControlNet and frozen Stable Diffusion network to transfer styles on the depth images. In the textual branch, DiffCLIP uses an optional Style-Prompt Generation Module for few-shot tasks and manual prompts for zero-shot tasks. Frozen CLIP image encoder and text encoder are used to generate feature representations of images and text which then go through a Multi-Modal Fusion Block.



Figure 2. **Style Transfer using Stable Diffusion and ControlNet:** Illustrating Results for 10 Categories (right) in the ModelNet10 Dataset Using a “Monitor” Depth Map (left). For example, when transferring the style of “monitor” to “bathtub”, the base of the monitor will be filled with blue rippling water. When transferring the style of “monitor” to “chair” or “sofa”, characteristic textures of these objects are displayed. When using the “monitor” label to transfer its own style, the resulting image clearly generates the toolbar and menu bar on the computer screen.

3.2.2 Stable-Diffusion-Based Style Transfer

To better use stable diffusion [22] in our task-specific domain, where the input is depth maps, we use ControlNet [37], a robust neural network training method to avoid overfitting and to preserve generalization ability when large models are trained for specific problems. In DiffCLIP, ControlNet employs a U-Net architecture identical to the one used in

stable diffusion. It duplicates the weights of an extensive diffusion model into two versions: one that is trainable and another that remains fixed. The linked trainable and fixed neural network units incorporate a specialized convolution layer known as “zero convolution.” In this arrangement, the weights of the convolution evolve from zero to optimized parameters in a learned manner. ControlNet has several implementations with different image-based conditions to control large diffusion models, including Canny edge, Hough line, HED boundary, human pose, depth, etc. We use the frozen pre-trained parameters of ControlNet under depth condition, which is pre-trained on 3M depth-image-caption pairs from the internet. The pre-trained ControlNet is then used to generate our own realistic images of depth maps using different class labels as prompts. We illustrate the style-transfer effects of ControlNet in Fig. 2.

3.2.3 Style-Prompt Generation

Optionally, when we want to do few-shot learning, we set up a style prompt generation module to embed style into prompts as is shown in the left part of Fig. 3.

We make use of “style features” in point cloud processing analogous to the usage in 2D vision. In 2D vision tasks, ConvNets act as visual encoders that extract features at dif-

ferent layers: higher-level layers tend to learn more abstract and global features while lower-level layers detect simple and local features such as edges, corners, and blobs [9; 35]. We conjecture that in the 3D point cloud case, the combination of features extracted from multiple levels could also improve the generalizability of the model, similar to how style features improve learning on RGB images in 2D vision tasks.

DiffCLIP leverages CLIP’s frozen vision encoder f_v and text encoder f_t . Additionally, a pre-trained point-based network, Point Transformer f_p [40], is used in DiffCLIP. The Point Transformer operates by first encoding each point in the point cloud into a high-dimensional feature vector using a multi-layer perception network [30], and then processing the encoded feature vectors are then processed by a series of self-attention layers, similar to the ones used in the Transformer [29] architecture.

We pre-train the Point Transformer network using our customized dataset, ShapeNet37, which consists of sampled data from ShapeNetCore [3]. The network comprises several blocks, each generating feature maps at different scales. To incorporate domain-specific characteristics, we project the output features from each block into a token and then concatenate them together with a class label to generate the style prompt as the input of CLIP’s text encoder.

3.3. Using DiffCLIP

3.3.1 Zero-shot 3D classification

For each 3D point cloud input x , after being projected into M views and densified, we obtain their smooth depth maps $x'_i, \forall i \in \{1, \dots, M\}$. These depth maps are then fed into ControlNet as conditions to guide the image generation of stable diffusion. We use the prompt “a photo of a [class], best quality, extremely detailed” as the default style prompt input to the stable diffusion model, denoted as t_j , where $j \in \{1, \dots, K\}$. The generated realistic images are represented as $x''_{ij} = d(x'_i, t_j)$, where the function $d(\cdot)$ denotes the stable diffusion combined with ControlNet. Based on the realistic images generated by the projection and style transfer stages, we employ CLIP to extract their visual features $f_{ij} = f_v(x''_{ij})$. In the textual branch, we use “a photo of a [class]” as the prompt and encode their textual features as the zero-shot classifier $W_t \in \mathbb{R}^{K \times D}$. Furthermore, the classification logits s_{ij} for each view i and each style transfer text guidance j are calculated separately as follows:

$$\text{logits}_{s_{ij}} = f_{ij} W_t^T, \forall i \in \{1, \dots, M\}, j \in \{1, \dots, K\} \quad (2)$$

For each style transfer text guidance j , the classification $\text{logits}_j \in \mathbb{R}^{1 \times K}$ is defined as

$\text{Max}P([\text{logits}_{s_{1j}}; \dots; \text{logits}_{s_{Mj}}])$, where $\text{Max}P(\cdot)$ represents taking the maximum value of each column in the matrix. Then, the probability matrix $P \in \mathbb{R}^{K \times K}$ is generated by $P = [\text{softmax}(\text{logits}_{s_1}); \dots; \text{softmax}(\text{logits}_{s_K})]$.

We design two strategies to calculate the final probability vector $\mathbf{p} = [p_1, p_2, \dots, p_K] \in \mathbb{R}^K$ for each classification. Each of them includes a global logits part and a local one in order to make full use of raw probability distributions of each diffusion result from each view on all the classes. The first strategy is given by

$$\mathbf{p} = \beta_1 \mathbf{p}_{\text{glo}} + \beta_2 \mathbf{p}_{\text{loc}} \quad (3)$$

where

$$\mathbf{p}_{\text{glo}} = \left[\sum_{i=1}^K P_{i1} \mathbb{I}(P_{i1} \leq P_{11}), \dots, \sum_{i=1}^K P_{iK} \mathbb{I}(P_{iK} \leq P_{KK}) \right] \quad (4)$$

$$\mathbf{p}_{\text{loc}} = [P_{11}, \dots, P_{KK}] \quad (5)$$

and β_1 and β_2 are hyper-parameters. \mathbf{p}_{glo} represents summing all values in the matrix that are no more than the diagonal by column, which represents the global information, and \mathbf{p}_{loc} returns the diagonal entries of P that represent the probabilities of the realistic images generated by text guidance j being classified into category j , which provides local information. We illustrate the detailed computation in our supplementary material. The second strategy to aggregate the matrix P to calculate the final logits is as follows:

$$\mathbf{p} = \text{norm}(\mathbf{p}_{\text{glo}}) * \mathbf{p}_{\text{loc}} \quad (6)$$

where

$$\begin{aligned} \mathbf{p}_{\text{glo}} &= \left[\sqrt[\kappa]{\prod_{i=1}^K P_{i1}}, \dots, \sqrt[\kappa]{\prod_{i=1}^K P_{iK}} \right] \\ \mathbf{p}_{\text{loc}} &= \left[\max_{i \in \{1, 2, \dots, K\}} P_{i1}, \dots, \max_{i \in \{1, 2, \dots, K\}} P_{iK} \right] \quad (7) \\ \text{norm}(\mathbf{p}) &= (\mathbf{p} - \min(\mathbf{p})) / (\max(\mathbf{p}) - \min(\mathbf{p})) \end{aligned}$$

Each entry in \mathbf{p}_{glo} represents the global feature of each category, which is defined as the geometric mean of the probabilities of that category across all style transfer results. Each entry in \mathbf{p}_{loc} represents the local feature for each category and is obtained from the diffusion result that is most similar to that category itself. Detailed clarification of these equations are shown in the appendix. The experiment results demonstrate that these two calculation strategies have their own strengths and weaknesses on different datasets.

To better illustrate the calculation process of probability matrix P mentioned in equation 3 and equation 5, we give the matrix in Fig. 4 as an example. In equation 6, while calculating \mathbf{p}_{glo} , numbers in blue boxes, which are bigger than numbers in orange boxes, are ignored.

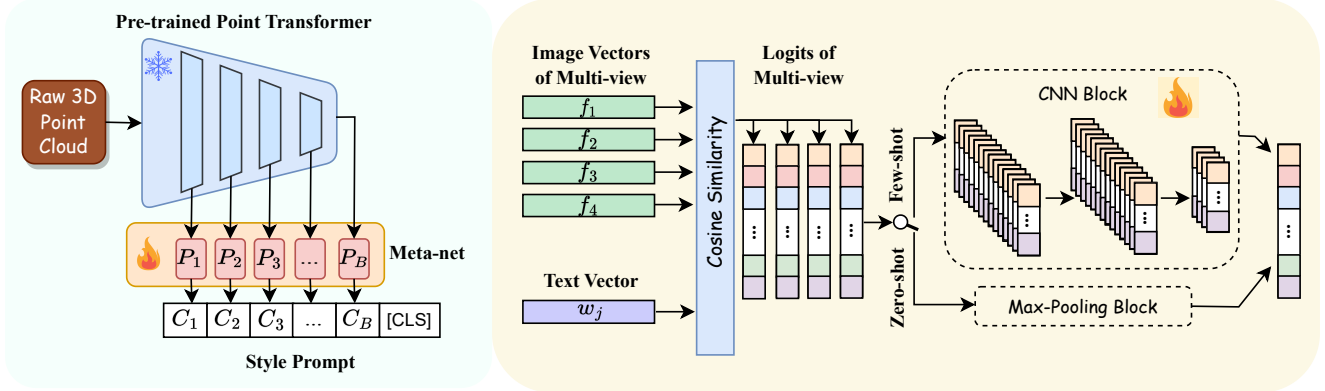


Figure 3. **Prompt generation module** (left) and **Multi-view Fusion Block** (right) of DiffCLIP.

		Logits Distribution on Target Classes									
		bathub	bed	chair	desk	dresser	monitor	nightstand	sofa	table	toilet
Diffusion Prompt	bathub	0.9120	0.0155	0.0143	0.0090	0.0052	0.0024	0.0184	0.0089	0.0061	0.0081
	bed	0.0225	0.8620	0.0088	0.0235	0.0103	0.0219	0.0109	0.0043	0.0113	0.0246
	chair	0.9520	0.0079	0.0051	0.0005	0.0036	0.0045	0.0055	0.0080	0.0024	0.0105
	desk	0.0749	0.0083	0.1210	0.5820	0.0375	0.0441	0.0183	0.0449	0.0218	0.0472
	dresser	0.0040	0.0005	0.0077	0.0049	0.9720	0.0032	0.0015	0.0003	0.0004	0.0066
	monitor	0.0967	0.0773	0.0136	0.0891	0.0456	0.5310	0.0349	0.0035	0.0112	0.0971
	nightstand	0.0008	0.0060	0.0015	0.0005	0.0018	0.5572	0.4310	0.0005	0.0002	0.0006
	sofa	0.0213	0.0522	0.0327	0.0346	0.0690	0.0394	0.0202	0.7230	0.0047	0.0029
	table	0.0001	0.0003	0.0129	0.0007	0.0056	0.0002	0.0101	0.0011	0.9670	0.0019
	toilet	0.0015	0.0002	0.0003	0.0010	0.0007	0.0005	0.0021	0.0019	0.0008	0.9910

Figure 4. An example of matrix P .

3.3.2 Few-Shot 3D Classification

In the context of few-shot classification tasks, a primary innovation of our model lies in the generation of style prompts. The image branch of the model remains identical to the one used in zero-shot classification tasks. As detailed in Section 3.2.3, we assume that the point-based network consists of B blocks. To incorporate domain-specific characteristics, we establish B meta-nets as style projectors $\mathcal{P}_b(\cdot; \theta_b)$, $\forall b \in \{1, \dots, B\}$ to encode domain characteristics into B prefix features c_b , where θ_b represents the parameters of the b -th meta-net. Assuming that the textual feature from the b -th block is D_b -dimension, the total dimension of our textual encoder is $D = \sum_{b=1}^B D_b$. Specifically, $c_b \in \mathbb{R}^{D_b}$ is computed by $c_b(x) = \mathcal{P}_b(\mathcal{F}_b(x); \theta_b)$.

To generate a style-prompt for a given (x, y) pair, where $x \in \mathbb{R}^{2048 \times 3}$ is an original 3D point cloud input, we define the full style prompt t_y by

$$t_y = [c_1(x)][c_2(x)] \dots [c_B(x)][CLS_y], \quad (8)$$

where $[CLS_y]$ is the word embedding of label y . Subsequently, a zero-shot classifier $W_t \in \mathbb{R}^{D \times K}$ can be generated.

Following the computation of logits, as in zero-shot clas-

Method	Zero-shot Performance				
	MN10	MN40	ScanObjectNN		
			OBJ_ONLY	OBJ_BG	PB_T50_RS
PointCLIP [39]	30.2	23.8	21.3	19.3	15.4
Cheraghian [43]	68.5	-	-	-	-
CLIP2Point [11]	66.6	49.4	35.5	30.5	23.3
PointMLP+CG3D [10]	64.1	50.4	-	-	25.0
PointTransformer+CG3D [10]	67.3	50.6	-	-	25.6
PointCLIP v2 [43]	73.1	64.2	50.1	41.2	35.4
ReCon [18]	75.6	61.7	43.7	40.4	30.5
DiffCLIP	82.4	54.2	45.3	43.2	35.2

Table 1. Zero-shot classification accuracy (%) of DiffCLIP on ModelNet10 [32], ModelNet40 [32], ScanObjectNN [28].

sification, we establish another trainable module, using a convolutional neural network to fuse logits s_i of each view $i \in \{1, \dots, M\}$, illustrated in the right portion of Fig. 3.

4. Experiments and Results

4.1. Zero-Shot Classification

We evaluate the zero-shot classification performance of DiffCLIP on three well-known datasets: ModelNet10 [32], ModelNet40 [32], and ScanObjectNN [28]. For each dataset, we require no training data and adopt the full test set for evaluation. For the pre-trained CLIP model, we adopt ResNet-50 [9] as the visual encoder and transformer [29] as the textual encoder by default, and then try several other encoders in the ablation studies. We initially set the prompt of stable diffusion as “a photo of a [class], best quality, extremely detailed”, and then we try multiple prompts to adapt the three different datasets later.

Performance. In Table 1, we present the performances of zero-shot DiffCLIP for three datasets with their best performance settings. Without any 3D training data, DiffCLIP achieves an accuracy of 43.2% on OBJ_BG of ScanObjectNN dataset, an accuracy of 82.4% for zero-shot classi-

	bathtub	bed	chair	desk	dresser	monitor	nightstand	sofa	table	toilet	Average
DiffCLIP	97.0	88.4	46.5	33.7	94.0	99.0	99.0	78.0	87.0	94.0	82.2
Ablation	83.0	62.8	3.5	29.1	36.0	100.0	98.0	98.0	89.0	80.0	71.4

Table 2. Ablation study. Line 2 shows zero-shot classification accuracy(%) of DiffCLIP without ControlNet and Stable diffusion module on ModelNet10.

fication on ModelNet10, which are state-of-the-art performance, and an accuracy of 35.2% for zero-shot classification on PB_T50_RS, which is comparable to state-of-the-art performance. While PointCLIP v2’s [43] results are sometimes better than ours, GPT-3 is used in the PointCLIP v2 model, which would require careful manual labor and further selection of the prompts that are useful for each sample in the test set. DiffCLIP does not require extensive manual tweaking of the prompts to achieve acceptable results by leveraging stable diffusion, which realizes full zero-shot classification without overfitting to each test set.

Ablation Studies. We conduct ablation studies on zero-shot DiffCLIP concerning the importance of stable diffusion plus ControlNet module, as well as the influence of different projection views and projection view numbers on ModelNet10. For zero-shot DiffCLIP structure without style transfer stages, in which the depth maps are directly sent into the frozen image encoder of CLIP, the best classification accuracy reaches 58.8%, while implementing ViT/16 as the image encoder of CLIP. Hence, we can see the style transfer stage increases the accuracy by 21.8%. Detailed ablation results are shown in Table 2.

In terms of the number of projected views, we try 1 (randomly selected), 1 (elaborately selected), 4, and 8 projection view(s) on ModelNet10. In Table 3, records are taken from a bird’s-eye view angle of 35 degrees, and the angles in the table represent the number of degrees of horizontal rotation around the Z-axis. The result of one elaborately selected view from a -135° angle outperforms others. Due to the characteristics of ModelNet10 dataset that all the objects are placed horizontally, projection from 4 views does not show obvious advantages. But the performance of the method on ModelNet40 is improved from multi-view projection because objects have different spatial arrangements, as shown in Table 5. We also ablate on the choice of different visual backbones, shown in Table 4. The results suggest that ViT/16 yields the best results in our framework.

Prompt Design. In DiffCLIP, we design several different prompts as both the textual branch of the CLIP model and the style prompt input of stable diffusion. For ModelNet10 and ModelNet40, we implement a default prompt for stable diffusion. The performance of multiple prompt designs for CLIP on ModelNet10 is shown in Table 6. We try to

place the class tag at the beginning, middle, and end of sentences, among which class tag located at the middle of the sentence performs best. Specifically, real-world objects in ScanObjectNN are incomplete and unclear, we elaborately design its diffusion style prompt as “a photo of a [class], behind the building, best quality, extremely detailed”, in order to generate an obstacle in front of the target object, which is shown to be more recognizable for CLIP’s image encoder.

Different Projection View Performance						
View	1(-135°)	1(-45°)	1(45°)	1(135°)	1(random)	4
ModelNet10	82.4	67.5	55.2	77.9	54.3	80.1

Table 3. Zero-shot classification accuracy (%) of DiffCLIP on ModelNet10 under various projection view(s) and view numbers, using ViT/16 as the visual encoder of CLIP.

Different Visual Encoders					
Models	RN50	RN101	ViT/32	ViT/16	RN.×16
PointCLIP	20.2	17.0	16.9	21.3	23.8
DiffCLIP	68.5	73.1	77.3	82.4	74.4

Table 4. Zero-shot classification accuracy (%) of DiffCLIP on ModelNet10 implementing different encoders with best projection view(s).

Prompt for CLIP	Accuracy
"a photo of a + [C]"	80.1
"a rendered image of + [C]"	78.1
"a 3D rendered image of + [C]"	79.4
"[C] + with white context"	78.0
"[C] + with white background"	79.9
"a 3D image of a + [C] + with rendered background"	82.4

Table 6. Zero-shot classification accuracy (%) of DiffCLIP on ModelNet10 using different prompt for CLIP.

	Single-view	Multi-view
ModelNet40	49.7	54.2

Table 5. Zero-shot classification accuracy(%) of DiffCLIP on ModelNet40 under different view numbers.

Method	ModelNet10			
	1-shot	2-shot	4-shot	8-shot
PointNet [16]	42.9	44.3	61.4	78.0
PointNet++ [17]	53.2	69.3	73.9	83.8
DiffCLIP	45.8	48.2	62.6	79.0

Table 7. Few-shot classification accuracy (%) of DiffCLIP on ModelNet10 with different shot numbers.

4.2. Few-shot Experiments

We experiment DiffCLIP with the trainable Style-Prompt Generation Module and the trainable Multi-view Fusion Module under 1, 2, 4, and 8 shots on ModelNet10. For K -shot settings, we randomly sample K point clouds from each category of the training set. We set the learning rate as $1e-4$ initially to train the meta-nets in the style prompt generation module, and then turn the learning rate of the meta-net into $1e-7$ and train the multi-view fusion block. Due to the limitation of memory of our GPU, the training batch size is set as 1. Default prompts are used for stable diffusion.

Performance. As is shown in Table 7, we compare few-shot classification performance of DiffCLIP to PointNet [16] and PointNet++ [17] in terms of K -shot classification accuracy. DiffCLIP surpasses PointNet in performance and approximately approaches PointNet++. It reaches an accuracy of 79.0% under 8-shot training.

5. Discussion and Limitations

In zero-shot 3D classification tasks, DiffCLIP achieves state-of-the-art performance on OBJ_BG of ScanObjectNN dataset and accuracy on ModelNet10 which is comparable to state-of-the-art. In few-shot 3D classification tasks, DiffCLIP surpasses PointNet in performance and approximately approaches PointNet++. However, our model does have some notable room for improvement. First, DiffCLIP’s performance on zero-shot tasks needs more exploration, such as projecting 3D point cloud into more views, designing better functions to calculate the logits vector \mathbf{p} . Second, DiffCLIP’s performance on few-shot tasks are still not that good which needs further improvement, such as fine-tuning on initial CLIP encoders as well as ControlNet. Moreover, the ability of DiffCLIP processing some other 3D tasks including segmentation and object detection is expected to be explored. We plan to leave these for future work.

Limitations. First, the implementation of stable diffusion introduces a time-intensive aspect to both the training and testing of the model due to the intricate computations required in the diffusion process, which can elongate the time for model execution 8. We also acknowledge that the scale of our pre-training dataset for the point transformer is presently limited. This constraint impacts the performance

Step for diffusion	Time consumption	Accuracy
0	94min	71.7
5	304min	80.7
10	549min	82.4
15	862min	82.2
20	1061min	80.6

Table 8. Time consumption of zero-shot inference and accuracy(%) of DiffCLIP with various diffusion steps on ModelNet10.

of DiffCLIP. A larger, more diverse dataset would inherently provide a richer source of learning for the model, thereby enhancing its capability to understand more data.

6. Conclusion

In conclusion, we propose a new pre-training framework called DiffCLIP that addresses the domain gap in 3D point cloud processing tasks by incorporating stable diffusion with ControlNet in the visual branch and introducing a style-prompt generation module for few-shot tasks in the textual branch. The experiments conducted on ModelNet10, ModelNet40, and ScanObjectNN datasets demonstrate that DiffCLIP has strong abilities for zero-shot 3D understanding. The results show that DiffCLIP achieves state-of-the-art performance with an accuracy of 43.2% for zero-shot classification on OBJ_BG of ScanObjectNN dataset and a comparable accuracy with state-of-the-art of 82.4% for zero-shot classification on ModelNet10. These findings suggest that the proposed framework can effectively minimize the domain gap and improve the performance of large pre-trained models on 3D point cloud processing tasks.

References

- [1] Shirsha Bose, Enrico Fini, Ankit Jha, Mainak Singha, Biplob Banerjee, and Elisa Ricci. Stylip: Multi-scale style-conditioned prompt learning for clip-based domain generalization. *arXiv preprint arXiv:2302.09251*, 2023. **2**
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. **3**
- [3] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. **5**
- [4] Hsin-Yu Chang, Zhixiang Wang, and Yung-Yu Chuang. Domain-specific mappings for generative adversarial style transfer. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII 16*, pages 573–589. Springer, 2020. **2**
- [5] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. **3**
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. **1**
- [7] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. **2**
- [8] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7:187–199, 2021. **2**
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **5, 6**
- [10] Deepti Hegde, Jeya Maria Jose Valanarasu, and Vishal M Patel. Clip goes 3d: Leveraging prompt tuning for language grounded 3d recognition. *arXiv preprint arXiv:2303.11313*, 2023. **1, 2, 6**
- [11] Tianyu Huang, Bowen Dong, Yunhan Yang, Xiaoshui Huang, Rynson WH Lau, Wanli Ouyang, and Wangmeng Zuo. Clip2point: Transfer clip to point cloud classification with image-depth pre-training. *arXiv preprint arXiv:2210.01055*, 2022. **1, 2, 6**
- [12] Bu Jin, Beiwen Tian, Hao Zhao, and Guyue Zhou. Language-guided semantic style transfer of 3d indoor scenes. In *Proceedings of the 1st Workshop on Photorealistic Image and Environment Synthesis for Multimedia Experiments*, pages 11–17, 2022. **2**
- [13] Lei Li, Siyu Zhu, Hongbo Fu, Ping Tan, and Chiew-Lan Tai. End-to-end learning local multi-view descriptors for 3d point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1919–1928, 2020. **2**
- [14] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 922–928. IEEE, 2015. **2**
- [15] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. **3**
- [16] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. **2, 8**
- [17] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. **2, 8**
- [18] Zekun Qi, Runpei Dong, Guofan Fan, Zheng Ge, Xiangyu Zhang, Kaisheng Ma, and Li Yi. Contrast with reconstruct: Contrastive 3d representation learning guided by generative pretraining. *arXiv preprint arXiv:2302.02318*, 2023. **6**
- [19] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. **1, 2**
- [20] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. **2**
- [21] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3577–3586, 2017. **2**
- [22] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. **1, 3, 4**
- [23] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. **3**
- [24] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, pages 36479–36494, 2022. **3**
- [25] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time

- prompt tuning for zero-shot generalization in vision-language models. *arXiv preprint arXiv:2209.07511*, 2022. [2](#)
- [26] Yang Shu, Xingzhuo Guo, Jialong Wu, Ximei Wang, Jianmin Wang, and Mingsheng Long. Clipood: Generalizing clip to out-of-distributions. *arXiv preprint arXiv:2302.00864*, 2023. [2](#)
- [27] Hang Su, Subhansu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 945–953, 2015. [2](#)
- [28] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit 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*, pages 1588–1597, 2019. [6](#)
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [5](#), [6](#)
- [30] Paul Werbos. Beyond regression: New tools for prediction and analysis in the behavioral sciences. *PhD thesis, Committee on Applied Mathematics, Harvard University, Cambridge, MA*, 1974. [5](#)
- [31] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7959–7971, 2022. [2](#)
- [32] 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 the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. [6](#)
- [33] Haoxuan You, Yifan Feng, Rongrong Ji, and Yue Gao. Pvnnet: A joint convolutional network of point cloud and multi-view for 3d shape recognition. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1310–1318, 2018. [2](#)
- [34] Tan Yu, Jingjing Meng, and Junsong Yuan. Multi-view harmonized bilinear network for 3d object recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 186–194, 2018. [2](#)
- [35] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, pages 818–833. Springer, 2014. [5](#)
- [36] Yihan Zeng, Chenhan Jiang, Jiageng Mao, Jianhua Han, Chaoqiang Ye, Qingqiu Huang, Dit-Yan Yeung, Zhen Yang, Xiaodan Liang, and Hang Xu. Clip²: Contrastive language-image-point pretraining from real-world point cloud data. *arXiv preprint arXiv:2303.12417*, 2023. [1](#), [2](#)
- [37] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023. [2](#), [3](#), [4](#)
- [38] Renrui Zhang, Rongyao Fang, Wei Zhang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free clip-adapter for better vision-language modeling. *arXiv preprint arXiv:2111.03930*, 2021. [2](#)
- [39] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Point-clip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022. [1](#), [2](#), [3](#), [6](#)
- [40] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021. [2](#), [5](#)
- [41] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. [2](#)
- [42] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [2](#)
- [43] Xiangyang Zhu, Renrui Zhang, Bowei He, Ziyao Zeng, Shanghang Zhang, and Peng Gao. Pointclip v2: Adapting clip for powerful 3d open-world learning. *arXiv preprint arXiv:2211.11682*, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)