

CLIP goes 3D: Leveraging Prompt Tuning for Language Grounded 3D Recognition

Deepti Hegde*, Jeya Maria Jose Valanarasu*, Vishal M. Patel
Johns Hopkins University

dhegde1, jvalana1, vpatel136@jhu.edu

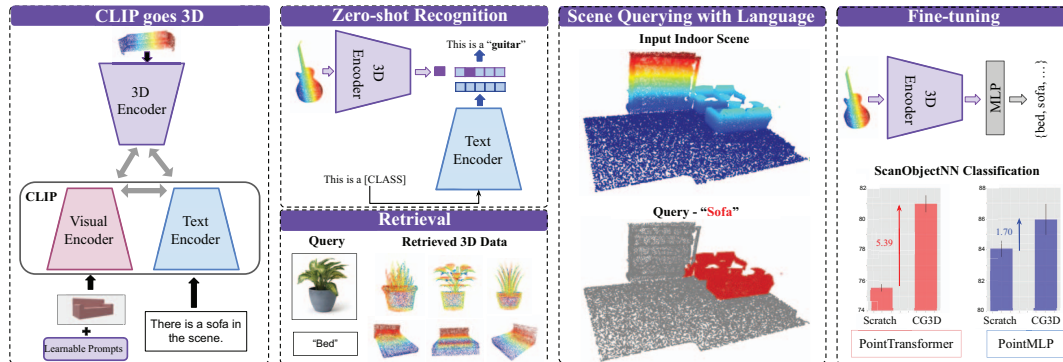


Figure 1: Overview of our proposed framework CLIP goes 3D (CG3D). We introduce a 3D Encoder in the CLIP framework and pre-train it using natural language supervision while also leveraging CLIP’s pre-trained visual encoder. CG3D solves various practical tasks like zero-shot 3D recognition, 3D point cloud retrieval, scene querying with natural language. Moreover, it can serve as a strong initial weight for standard fine-tuning tasks.

Abstract

Vision-Language models like CLIP have been widely adopted for various tasks due to their impressive zero-shot capabilities. However, CLIP is not suitable for extracting 3D geometric features as it was trained on only images and text by natural language supervision. We work on addressing this limitation and propose a new framework termed CG3D (CLIP Goes 3D) where a 3D encoder is learned to exhibit zero-shot capabilities. CG3D is trained using triplets of pointclouds, corresponding rendered 2D images, and texts using natural language supervision. To align the features in a multimodal embedding space, we utilize contrastive loss on 3D features obtained from the 3D encoder, as well as visual and text features extracted from CLIP. We note that the natural images used to train CLIP and the rendered 2D images in CG3D have a distribution shift. Attempting to train the visual and text encoder to account for this shift results in catastrophic forgetting and a notable decrease in performance. To solve this, we employ prompt tuning and introduce trainable parameters in the input space to shift CLIP towards the 3D

pre-training dataset utilized in CG3D. We extensively test our pre-trained CG3D framework and demonstrate its impressive capabilities in zero-shot, open scene understanding, and retrieval tasks. Further, it also serves as strong starting weights for fine-tuning in downstream 3D recognition tasks. Code: <https://github.com/deeptihgde/CLIP-goes-3D>

1. Introduction

For many tasks in 2D vision, the most efficient and accurate results are now obtained by adapting foundation models [42, 43, 2, 64] which are pre-trained on large-scale data. There is currently a significant amount of research focused on efficiently adapting foundation models for specific 2D vision tasks [65, 47, 59, 15], rather than developing new supervised methods from scratch. These approaches seek to build on existing models and leverage their pre-trained features to achieve better performance on target tasks with less data and computation requirements. These recent trends in vision are actually similar to trends that were observed in natural language processing (NLP) a few years ago. In NLP, foundation models have been dominating since 2018 as models like BERT [9] and GPT-3 [4] showed excep-

*Equal Contribution

tional ability to accomplish various NLP tasks, such as question answering, sentence prediction, sentiment classification, etc. Also, foundation models pre-trained on multimodal data like images and text have been useful by exhibiting impressive zero-shot capabilities. In particular, Contrastive Language Image Pre-training (CLIP) [41] has been applied to various 2D tasks, including image classification [71], object detection [48, 13], image segmentation [25, 57, 69], image retrieval [18, 29], and visual question answering [34, 8]. In this work, we try to bridge the gap between 2D vision language models and 3D shape understanding and focus on answering the following question: *How can we build a 3D network that can possess similar functionalities as that of a foundation model like CLIP?*

3D visual understanding has many practical applications in robotics [16, 70], augmented reality [17, 55, 62], and autonomous driving [39, 38, 46, 33]. Comprehending the semantics and characteristics of each point in a 3D space is crucial for addressing a wide range of issues in downstream tasks. Thus, there lies several use-cases for a potential 3D foundational model similar to foundation models for vision and NLP. A powerful 3D network with zero-shot capabilities does not only help improve the performance of existing 3D backbones but also enables open 3D scene understanding and 3D retrieval tasks. However, the development of foundational models for 3D understanding faces several challenges, including the limited availability of 3D data compared to images. While CLIP was able to leverage the vast amount of images available on the internet to create a large pre-training dataset of image-caption pairs, a similar approach cannot be directly applied to pre-train a 3D encoder with text due to the scarcity of 3D data. There have been some recent works like PointCLIP [66] which tried to utilize CLIP’s zero-shot capabilities for 3D zero-shot problems. PointCLIP directly uses the depth maps of a 3D point cloud on 2D visual encoder of CLIP. While it provides a quick and simple solution for 3D zero-shot problems, it lacks the characteristics of a foundational model since it cannot be used for 3D fine-tuning tasks or for 3D open scene understanding. Furthermore, it does not possess the ability to extract any 3D geometric features relevant for downstream tasks in 3D understanding.

To this end, we propose a new pre-training framework termed CG3D (CLIP Goes 3D) that trains a 3D encoder using natural language supervision while also leveraging CLIP’s knowledge. We begin by creating a pre-training dataset consisting of triplets of 3D point clouds, images, and corresponding text descriptions. We use point clouds from ShapeNet [5] as our 3D data and curate its corresponding rendered 2D image and a caption. Since ShapeNet consists of textured CAD models, we render random views of each object to use as the image pair. Despite their distinct properties, both a 3D point cloud and an image of the same ob-

ject share common semantic attributes. This is affirmed by the success of tasks such the transfer of pre-trained weights from an image-based network to a 3D point cloud classification network [60]. CG3D aims to ensure that there is similarity between 3D features and 2D features, as well as between 3D features and text features for objects of the same category, while being dissimilar for objects of different categories. This contrastive approach to learning enables the 3D encoder to acquire zero-shot capabilities similar to those of CLIP.

The process of training the 3D encoder with contrastive loss and comparing 3D features to the 2D features from CLIP’s visual encoder is a means of distilling CLIP’s semantic features to the 3D encoder. Although it would be efficient to train CLIP’s visual encoder to align with the data distribution of 3D objects and their related images, we observed a significant decrease in performance when training both CLIP’s visual and text encoders along with the 3D encoder in CG3D. This can be explained as CLIP starts to catastrophically forget its previous features while being trained to shift to the new distribution. However, it is not ideal to keep the visual encoder completely frozen. Large-scale language models like CLIP are trained mostly on natural images, which differ in distribution to the graphically rendered views of 3D objects. To address this domain gap [53, 51, 52, 50], we propose using prompt tuning techniques [19, 7] to shift the distribution in the input space before forwarding it to the visual encoder. We add visual prompts to the transformer backbone of CLIP’s visual encoder thus adding only a small amount of parameters in the input space while keeping the weights of visual encoder of CLIP frozen. These parameters learn the shift in input distribution to suit CLIP so that the 3D pre-training is effective. To demonstrate the effectiveness of CG3D, we conduct several experiments. First, we show its zero-shot capabilities on synthetic and real object datasets like ModelNet [56] and ScanObjectNN [49]. Additionally, we showcase the 3D model’s ability in open-scene comprehension by utilizing text-based queries, as well as its ability to conduct cross-modal 3D data retrieval while utilizing image or text queries. Further, the weights obtained from pre-training the 3D encoder using CG3D can also serve as effective initial weights when fine-tuning the model for other 3D tasks.

In summary, the following are our major contributions:

- We propose CG3D, a contrastive pre-training framework for training 3D networks using natural language supervision while also leveraging the knowledge of CLIP.
- We utilize prompt tuning to shift the input space of a pre-trained visual encoder from rendered images of CAD objects to natural images, allowing for more effective use of CLIP for 3D shapes.
- We conduct extensive experiments to demonstrate the versatile capabilities of CG3D. It exhibits strong zero-shot,

3D retrieval and 3D scene understanding capabilities with language. CG3D also acts as strong starting weights for multiple 3D recognition tasks.

2. Related works

Vision Language Models: The use of large-scale text-pre-training on attention-based models [9, 67] has led to the increasing popularity of vision-language models (VLM) due to their impressive performance in visual understanding tasks [22, 28, 41]. Recent advancements in contrastive learning have enabled CLIP [41] to perform multimodal learning with 400M noisy data crawled from the web. CLIP has been extended for high efficiency model training and cycle consistency through various methods, such as ALBEF [20] and Cyclip [12]. Adopting VLM for 3D point cloud processing is still in its infancy. PointCLIP [66] was the first method to adopt CLIP for 3D tasks. It directly uses the depth maps of 3D point clouds and uses it on the visual encoder of CLIP to perform zero-shot classification. Unlike this, we focus on using a 3D encoder in the CLIP so that it can directly take in a 3D point cloud.

3D Point cloud processing methods: In general, point cloud processing methods either process the original point cloud sets directly [39, 40] or transform the original point clouds into intermediate representations such as voxels [31, 45] or images [61, 21]. PointNet [39] enabled the direct use of unordered point sets as input through shared MLPs. PointNet++ [40] was later proposed as an extension of PointNet, incorporating a hierarchical feature learning approach that recursively captures local geometric structures. Recently, methods like PointTransformer [68] and PCT [14] have proposed transformer-based methods showing a significant improvements in performance. The current state of the art method is PointMLP [30] which effectively uses a deep residual MLP network for point cloud analysis.

There have also been pre-training methods which show 3D backbones can be pre-trained with unlabelled data to obtain strong initial weights to boost fine-tuning performance. PointContrast [58] performs contrastive training [6] by pushing together heavily augmented views of the same sample and minimizing the similarity between views of other samples on point cloud scenes that have undergone rigid transformations. CrossPoint [1] boosts point cloud classification performance by maximising the agreement between images and point cloud objects. PointBERT [63] and PointMAE [37] leverage masked modelling methods to perform pre-training. Unlike these backbone or pre-training works, we focus on enabling zero-shot capabilities to given 3D encoder in our proposed CG3D framework.

3. Method

Our main objective is to train a 3D shape encoder to acquire shape characteristics that can effectively capture the




Figure 2: Overview of the proposed learning strategy in CG3D. Note that only the 3D Encoder and learnable visual prompts are trained while everything else is frozen.

geometric properties of point clouds while also aligning with CLIP’s feature representation for each object category. In essence, we aspire to acquire features that are unique to each category yet unaffected by the mode of representation. To this end, we use point cloud-image-caption triplets to train the framework. Each element within a triplet of point cloud-image-caption is indicative of an object possessing specific semantic traits that are shared among the other objects in the collection. In this section, we first give an overview of the proposed CG3D framework, followed by details to effectively train the network. We then give details on the potential use-cases of CG3D.

3.1. CG3D framework

The CG3D framework as illustrated in Fig 2 consists of 3 networks - 3D shape encoder, visual encoder and text encoder from CLIP.

3D Encoder takes in a 3D point cloud as the input. To capture the essential shape characteristics of an object, we employ a 3D encoder that is specifically designed to analyze point clouds and generate a feature vector representing the object. Our framework is agnostic with the choice of 3D encoder, and an added projection layer ensures the output feature dimension remains consistent.

2D Visual Encoder takes in the corresponding rendered image of the 3D pointcloud as the input. Although shape features are essential in representing point clouds, image backbones from vision-language models trained on large amounts of data offer powerful feature representations of images that are semantically correlated with text. We employ CLIP’s visual encoder as it is pre-trained on a massive

amount of data. By utilizing the visual encoder, we can acquire highly effective and implied representations of categories present in point cloud datasets, which are then used to align with the 3D features. Note that CLIP provides both ResNet and ViT weights for visual encoder.

Text Encoder takes in the corresponding text caption of the 3D point cloud as input. Adopting natural language supervision for feature learning has been successful in training models to grasp visual concepts that can be depicted in both images and text [41]. We use the CLIP text encoder out-of-the-box under the same configuration as [41].

3.2. Training the 3D encoder

Our main training objective is to align the 3D point clouds with their corresponding category level images and texts. This alignment happens in a common embedding space to which data from each modality is projected from the modality-specific encoder and a projection head. Our training strategy relies on contrastive learning, which incentivizes cross-modal features of the same pair to be in close proximity to one another in the embedding space while keeping apart samples belonging to other pairs. We formulate the proposed losses below. Consider a set of N pointcloud-image-text triplets $\{x_i^{3D}, x_i^{2D}, x_i^{text}\}_{i=1}^N$, where x_i^{3D} represents a pointcloud, x_i^{2D} is the corresponding rendered image, and x_i^{text} is the corresponding text. Let the encoder for each modality be represented as ϕ_{3D}, ϕ_{2D} , and ϕ_{text} . We obtain the feature representation of each sample in a common embedding space by projecting the encoded feature to a common dimension represented by:

$$f_i^{3D} = \psi_{3D}(\phi_{3D}(x_i^{3D})) \quad (1)$$

$$f_i^{2D} = \psi_{2D}(\phi_{2D}(x_i^{2D})) \quad (2)$$

$$f_i^{text} = \psi_{text}(\phi_{text}(x_i^{text})) \quad (3)$$

where i ranges from 1 to the number of samples N and ψ is the projection operation for each modality. Through normalization, we constrain the output of each projection network to reside within a unit hypersphere, enabling us to measure feature similarity using the inner product. The contrastive losses (InfoNCE objective [36]) between 3D-image features; and 3D-text features are defined by:

$$\mathcal{L}_{(3D,2D)} = \frac{1}{2} \sum_{(f^{3D}, f^{2D}) \in B} \text{NCE}(f^{3D}, f^{2D}) + \text{NCE}(f^{2D}, f^{3D}) \quad (4)$$

$$\mathcal{L}_{(3D,text)} = \frac{1}{2} \sum_{(f^{3D}, f^{text}) \in B} \text{NCE}(f^{3D}, f^{text}) + \text{NCE}(f^{text}, f^{3D}) \quad (5)$$

where f^{3D} , f^{2D} , and f^{text} are the projected 3D, image and text features respectively. B corresponds to the batch. NCE loss is defined as:

$$\text{NCE}(f^A, f^B) = -\log \frac{\exp(\langle f^A, f_+^B \rangle / \tau)}{\sum_{f \in \{f_+^B, f_-^B\}} \exp(\langle f^A, f \rangle / \tau)} \quad (6)$$

where A, B are two different modalities and f^A, f^B are their corresponding features. τ is the temperature hyperparameter, f_+^B are the positive embeddings from modality B overlapping with modality A , and f_-^B are the negatively embeddings formed while pairing with modality A . For example, in $\mathcal{L}_{(3D,2D)}$ the positive pairs between 3D and 2D are formed by matching the features corresponding to same class while the rest are termed as negative pairs. The total loss \mathcal{L}_{3D} used to train the 3D encoder is defined by:

$$\mathcal{L}_{3D} = \mathcal{L}_{(3D,2D)} + \mathcal{L}_{(3D,text)} \quad (7)$$

3.3. Prompt tuning for visual encoder

The visual encoder in the CG3D framework takes in the rendered image of the 3D point cloud as its input. While the CLIP visual encoder has been trained on vast amounts of internet data and is highly resilient, during pre-training of CG3D, it only deals with rendered images. As a result, fine-tuning the CLIP visual encoder to handle rendered images could improve the training process for the 3D encoder. One possible approach is to train the visual encoder by optimizing its weights using the CLIP loss function that computes the similarity between image and text features. However, when we tried this method, we noticed a substantial decrease in performance. This phenomenon can be explained by the fact that training the visual encoder of CLIP causes catastrophic forgetting. That is, the encoder loses all of its prior knowledge while attempting to adapt to the new data distribution. Typically, this issue can be avoided by increasing the amount of new data available for fine-tuning. However, it's not feasible to obtain a large enough 3D dataset that matches the scale of the massive image-text data used to train CLIP. Therefore, we concentrate on developing methods to effectively fine-tune the model with new pre-training data while keeping the visual encoder frozen.

Visual prompt tuning, as described in [15], is a method that involves adding a small number of trainable parameters in the input space to fine-tune a base model for a specific task. In our proposed method for pre-training CG3D, we adopt this approach and modify the input space to better align with the visual encoder of the original CLIP model. This, in turn, allows the visual encoder to produce higher-quality features, which can enhance the training of the 3D encoder. We use deep prompting where we introduce learnable prompts as learnable tokens at every layer in the transformer layer in ViT (visual encoder). For an i^{th} transformer layer L_i , we define the collection of learnable tokens at that layer as $P_i = \{p_i^k, 1 \leq k \leq n\}$ where p corresponds to individual tokens and n is the total number of learnable tokens. The deep prompted visual encoder at i^{th} can be represented as:

$$[y_i] = L_i([y_{i-1}, P_{i-1}]) \quad (8)$$

where y_i is the output and y_{i-1} is the input to the current layer. Prompt tokens P are trained along with the 3D en-

coder in our CG3D framework. We use the original CLIP loss which is a contrastive loss between the image and text features to train these prompts. We formulate the loss used to train the prompts \mathcal{L}_P as:

$$\mathcal{L}_P = \frac{1}{2} \sum_{(f^{2D}, f^{text}) \in B} (\text{NCE}(f^{text}, f^{2D}) + \text{NCE}(f^{2D}, f^{text})) \quad (9)$$

where where f^{2D} and f^{text} are the image and text features respectively. B corresponds to the batch.

3.4. Using CG3D

3.4.1 Zero-shot 3D recognition

Zero-shot 3D classification refers to the method of classifying 3D objects without requiring any previous training on those specific objects. It has a number of use-cases in robotics and autonomous systems where objects need to be recognized quickly and accurately, without the need for extensive training on new objects. CG3D enables zero-shot 3D recognition directly with a 3D encoder which extracts shape features from the input point clouds. For zero-shot inference, we only use the 3D and text encoder from our proposed CG3D framework. The model takes as input a point cloud test sample, denoted as x_{test} , and a set of prompts $\mathbf{T} = \{t_i, 1 \leq i \leq N\}$ where t_i represents individual text prompts and N is the total number of text prompts. Each prompt t_i is denoted in the format of "This is a OBJECT", where OBJECT represents the name of a test class. The model's inference procedure is similar to that of CLIP, where it calculates the similarity score between each prompt and the test sample and selects the prompt that yields the highest score as the final prediction. This is formulated as:

$$y_{pred} = \max(\text{softmax}(\langle f^{3D}, \mathbf{F}^{text} \rangle)) \quad (10)$$

where f^{3D} and \mathbf{F}^{text} are the feature vectors of the point cloud and text inputs respectively, y_{pred} is the class prediction. Note that \mathbf{F}^{text} is actually a collection of feature vectors collected from forwarding the text queries T to the text encoder. This process has been summarized in Fig 1.

3.4.2 Scene querying with language

3D scene understanding is an important task in computer vision. It is crucial for enabling human-robot interaction and facilitating intuitive human-machine interfaces. In particular, querying a scene with language queries to understand the key details of a scene is a practically useful task. One aspect of this is answering queries such as "What is the location of the sofa?" or "Where can the chair be found?", which can help an individual in comprehending the environment or allow a robot to interact with it intelligently. By accurately identifying objects and their locations within a scene, robots can perform a wide range of tasks such as manipulation, navigation, and object recognition. Meanwhile,




Figure 3: Example of scene querying with text for a random indoor scene from S3DIS [3] dataset.

individuals can use interactive query and visualization tools to better understand and analyze complex 3D scenes.

CG3D enables zero-shot scene understanding with language queries. Without training the model on indoor scenes or using direct supervision, we show that a pre-trained CG3D framework can be used to understand scenes with text queries. Given an input scene, first we use k-means clustering to divide the scene into meaningful segments. We chose the best k value based on empirical analysis. Next, we feed forward all these clusters to the 3D encoder in our CG3D framework and get a set of 3D features $\mathbf{F}^{3D} = \{f_i^{3D}, 1 \leq i \leq k\}$. Here, k is the total number of clusters obtained from the scene and f_i^{3D} is the 3D feature vector from 3D encoder of CG3D for the i^{th} cluster. We also pass the text query to the text encoder to obtain f^{text} . Now, we match these 3D features with the text feature we obtain by forwarding the input query to the text encoder which can be denoted as:

$$y_{pred} = \max(\text{softmax}(\langle f^{text}, \mathbf{F}^{3D} \rangle)) \quad (11)$$

where y_{pred} is the predicted class. To demonstrate an example for scene understanding with language using CG3D, we pick random indoor scene as seen in Fig 3 from the S3DIS [3] dataset and use the query "Bookshelf". It can be observed that the model correctly classifies which cluster is the bookshelf in the scene. Although CG3D is not explicitly designed to provide a complete solution for open world 3D scene understanding, it facilitates scene querying using natural language.

3.4.3 Retrieval

3D point cloud retrieval is the process of searching and retrieving 3D point cloud data that is similar to a given query. The query here can be of different modality like an image or a text. Additionally, this can be reversed to retrieve images from a 3D query. This has several practical applications like matching a real-world object with its corresponding 3D

point cloud in a virtual environment to help create more realistic and accurate augmented and virtual reality experiences. To retrieve data using CG3D, we utilize the pre-trained encoders to obtain feature representations for both the query and the 3D point clouds. Specifically, we feed the image or text query into the corresponding encoder to obtain its feature vector, and we do the same for the 3D point clouds using the 3D encoder. The point clouds being forwarded to the 3D encoder constitute the complete database from which we are retrieving the relevant shapes. Next, we obtain the similarity score between these query feature and the 3D features and select the point clouds with the highest similarity scores as the output.

In Fig 4, we demonstrate the effectiveness of CG3D in retrieving relevant 3D point clouds and images. We randomly select some web images as queries, with ModelNet40 serving as our 3D database for shape retrieval. We pick the top four 3D point clouds that are of the highest similarity to the input query and display them in Fig 4. Similarly, the model is queried by text to retrieve shapes. Lastly, we perform retrieval in the reverse, by using point clouds as queries to obtain images from a small database of web-scraped images, experimenting with adding noise at the input to demonstrate robustness. It can be observed that all the retrieved point clouds and images are of very high similarity to the input query, despite the presence of noise and heavily occluded objects, proving the effectiveness of CG3D.

3.4.4 Fine-tuning for supervised tasks

Pre-training techniques are an effective strategy for enhancing the performance of fine-tuning in 3D computer vision tasks. Pre-training models on large datasets of unlabeled images can help them learn generic and transferable features, making them more robust to variations in data and enabling them to generalize well to new tasks and datasets.

Although the main objective of CG3D is its zero-shot capabilities, it also has the potential to serve as a valuable starting point for fine-tuning 3D models for downstream tasks. This is due to the excellent feature representation capabilities of the 3D encoder, which has been pre-trained using natural language supervision in the CG3D framework. Additionally, CG3D is model-agnostic, meaning that any 3D backbone can be pre-trained using CG3D, and the resulting weights can be used as a starting point for downstream tasks. We present multiple experiments that demonstrate the effectiveness of using CG3D for fine-tuning tasks.

4. Experiments and results

4.1. Datasets

Pre-training dataset: We choose ShapeNet [5] as the pre-training dataset due to its large number of classes and sam-

ples. ShapeNet consists of textured CAD models of 55 object categories and 52,460 total samples. We sample point clouds of a fixed size from each object mesh and normalize them to fit into a unit sphere. We render the colored CAD model views in Blender to obtain a image pair for each point cloud, following [11]. The text captions of each point cloud-image pair are framed as a descriptive sentences obtained from a set of standard templates such as “A photo of a {OBJECT}”. Each input point cloud is augmented with standard techniques such as object scaling, rotation, random drop, and perturbations.

Fine-tuning datasets: We perform zero shot (ZS) classification and downstream fine-tuning on the popular 3D datasets ModelNet40 [56] and ScanObjectNN [49]. As is standard, we evaluate on the full dataset (ModelNet40) as well as a 10-class subset (ModelNet10) for ZS classification. ModelNet40 consists of 12,311 synthetic meshes of common objects from 40 categories. Each mesh is down-sampled and normalized to fit a unit sphere. ScanObjectNN is a real-world point cloud dataset of objects from laser-scanned indoor scenes.

4.2. Implementation details

There exist several vision-language models that can be considered variants of CLIP that give superior zero-shot performance on images [20, 35, 24]. We leverage the pre-trained visual and text encoder weights from SLIP [35] to train CG3D, due to its performance and flexibility. We specifically choose ViT-Base [10] as the image backbone.

Pre-training: During pre-training, the visual prompt parameters and the parameters of the 3D encoder are tuned under different optimizers in alternate iterations. This is due to the fact that they are each supervised by disjoint loss functions and require different learning rates. We append 5 learnable prompt tokens at the input of every encoder layer in ViT, and initialize them randomly. The visual prompts are tuned using the SGD [44] optimizer under a cosine annealing [26] scheduler, with learning rate of 2×10^{-3} , weight decay 10^{-4} and an minimum learning rate of 10^{-6} . We follow the training convention of each 3D backbone. In case of PointTransformer [68] the network is tuned using the AdamW [27] optimizer under a cosine annealing [26] scheduler with a learning rate of 5×10^{-5} and a weight decay of 0.05, with a minimum learning rate of 10^{-6} . The PointMLP backbone is tuned under the same optimizer-scheduler scheme with learning rate 10^{-4} and weight decay 0.01. The entire framework is pre-trained for 100 epochs with a batch size of 32.

Fine-tuning: PointMLP is finetuned using the SGD optimizer and cosine scheduler with a learning rate of 0.02, a weight decay of 2×10^{-4} , and a minimum learning rate of 5×10^{-3} . PointTransformer is fine-tuned using the AdamW optimizer and cosine scheduler, with learning rate 2×10^{-4} ,

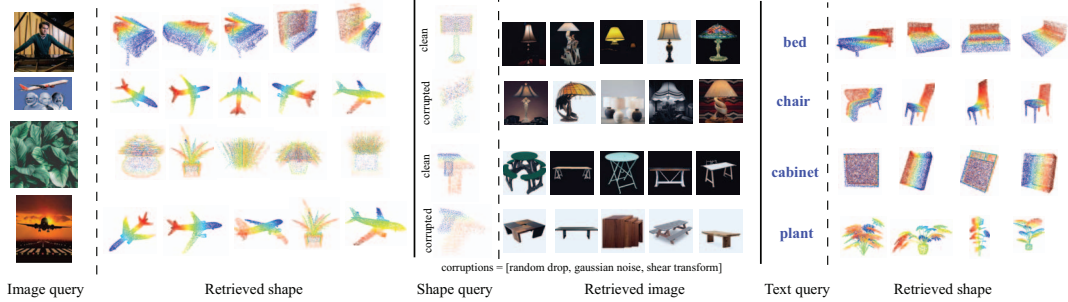


Figure 4: Retrieval results for CG3D.

Method	Zero-shot performance		
	MN10	MN40	ScanObjectNN
PointCLIP [66]	30.2	20.2	15.4
PointTransformer [68]+ CG3D	67.3	50.6	25.6
PointMLP [30] + CG3D	64.1	50.4	25.0

Table 1: Comparison of zero-shot classification performance of CG3D against that of PointCLIP for the ModelNet10, ModelNet40, and ScanObjectNN datasets.

a weight decay of 0.05 and a minimum learning rate of 10^{-6} . Each network is fine-tuned for 300 epochs with a batch size of 32 for PointMLP and 64 for PointTransformer. Our method is prototyped in PyTorch and all experiments are performed in a 8 GPU NVIDIA A100 cluster.

4.3. Zero-shot experiments

We present the results of zero-shot experiments conducted on test distributions of ModelNet10, ModelNet40, and ScanObjectNN in Table 1. We experiment with two backbones: PointTransformer and PointMLP pre-trained with CG3D. Note that the previous method PointCLIP uses a 2D depth map and CLIP’s visual encoder to get the prediction. We directly use the 3D encoder and extract relevant 3D shape features to perform the zero-shot classification. This gives a significant improvement over PointCLIP with an increase of 37.1% on ModelNet10, 30.4% on ModelNet40, and 10.2% on ScanObjectNN.

4.4. Fine-tuning Experiments

We present the results of our fine-tuning experiments in Table 2 on both synthetic (ModelNet40) and real (ScanObjectNN) datasets. For ScanObjectNN, we pick the hardest variant PB-T50-RS for our experiments. We compare against leading backbones as well as pre-training methods like Point-BERT and Point-MAE. We show fine-tuning performance on two backbones PointTransformer and PointMLP. It should be noted that our framework was not primarily developed to be a pre-training strategy, but rather to enable zero-shot capabilities for a 3D encoder.

Method	Overall accuracy	
	ModelNet40	ScanObjectNN
Pointnet [39]	89.2	68.0
Pointnet++ [40]	90.5	77.9
PointCNN [23]	92.2	78.5
DGCNN [54]	92.9	78.1
Point-BERT [63]	93.2	83.07
Point-MAE [37]	93.8	85.18
PointTransformer [68]	91.62 ± 0.29	75.56 ± 0.24
PointTransformer [68] + CG3D	92.93 ± 0.06	80.95 ± 0.54
PointMLP [30]	92.61 ± 0.13	84.08 ± 0.55
PointMLP [30] + CG3D	93.35 ± 0.18	85.78 ± 0.75

Table 2: Comparison of fine-tuning performance of CG3D with initial weights on ModelNet40 and ScanObjectNN (hardest variation: PB-T50-RS) against previous methods.

Even then, our framework demonstrates competitive performance as a pre-training strategy as observed in Table 2. In particular, we obtain a boost of 5.39% , 1.31% while using CG3D starting weights than random weights for PointTransformer on ScanObjectNN and ModelNet40 respectively. Also, we obtain a boost of 1.7% , 0.74% while using CG3D starting weights than random weights for PointMLP on ScanObjectNN and ModelNet40 respectively. These observations show the versatility of our proposed framework. We re-run the from-scratch experiments for the PointTransformer and PointMLP networks. To account for variability, we conducted the finetuning experiments for PointTransformer and PointMLP three times each, starting from scratch and also starting with CG3D weights. To account for this, we have reported the mean and standard deviation of the results obtained from these experiments.

5. Discussion

Ablation Study: We conduct an ablation study to analyze the role that each component of CG3D has on zero-shot performance. In Table 3, we report the overall zero-shot accuracy of the PointTransformer 3D encoder pre-trained under different loss configurations for the PB-T50-RS variant of

ScanObjectNN. We start with the configuration of training the 3D encoder with $\mathcal{L}_{(3D,2D)}$ and $\mathcal{L}_{(3D,text)}$ individually. We note that training with $\mathcal{L}_{(3D,text)}$ gives a slight improvement over just training with $\mathcal{L}_{(3D,2D)}$. Next, we pre-train CG3D with both these losses which obtains better performance over individual configurations. After this, we incorporated visual prompts into the CG3D model and trained it using \mathcal{L}_P . This further improved the model’s performance.

$\mathcal{L}_{(3D,2D)}$	$\mathcal{L}_{(3D,text)}$	visual prompt	ZS
✓	×	×	19.1
×	✓	×	19.7
✓	✓	×	23.9
✓	✓	✓	25.6

Table 3: Ablation Study on ScanObjectNN.

Analysis on Less Data: Pre-trained models are particularly useful when dealing with tasks that have limited access to data. To prove the effectiveness of CG3D in such scenarios, we conduct experiments on ModelNet40 with both PointMLP and PointTransformer as backbones. We fine-tune each 3D backbone on different sub-sets 10%, 20%, 30%, 50% of the data and present the results in Fig. 5. It can be observed that the model trained with starting weights of CG3D always obtains better performance





Figure 5: Experiments on data scarce setups on ModelNet40 with PointTransformer and PointMLP backbones.

Visualization of 3D Feature Representations: For analysis of the quality of features learned by our method, we visualize the UMAP [32] embeddings of 3D and image features extracted by CG3D while using ModelNet10. In Fig. 6, we plot the 3D feature learned by PointTransformer before and after pre-training with CG3D. Most features lack class separability before pre-training as can be seen in Fig. 6a. After pre-training, the 3D encoder is able to produce class discriminative features even for unseen categories in ModelNet10 as seen in Fig. 6b.


Effect of Prompt Tuning: We visualize the image features learned by the visual encoder with and without the learned prompts in CG3D. Since textured CAD models are not available for all samples, we consider depth maps of the points projected in the 2D plane. Figure 7a shows the depth map image features learned by the CLIP visual encoder. Since these images are visually dissimilar from natural images, the encoder fails to produce discriminative features. However, as seen in Figure 7b, the visual encoder



(a) Before CG3D. (b) After CG3D.

Figure 6: Comparison of UMAP embeddings of point cloud features from 3D encoder of CG3D while using ModelNet10 3D point clouds.

trained with visual prompts after CG3D pre-training produces features with improved class separability proving the effectiveness prompt tuning.



(a) Without VPT. (b) After contrastive VPT.

Figure 7: Comparison of UMAP embeddings of visual encoder features of CG3D with ModelNet10 depth maps.

Limitations: We note that our pre-training dataset is still small in size, and consists of only simulated point cloud objects, thus limiting the potential of CG3D. To build a powerful foundation model for 3D, we need to work on data curation of 3D pointclouds, with corresponding images and text captions. We also focused on pre-training on objects and not scenes. Pre-training on scenes could open up interesting full scene understanding capabilities in CG3D.

6. Conclusion

In this paper, we proposed a new framework CG3D, where a 3D Encoder is introduced into the CLIP framework. This 3D Encoder is trained such that the extracted 3D features align with the image and text features of the same category. We also proposed using learnable visual prompts to shift the rendered image distribution to that of the CLIP to get better representative image features from the visual encoder. Through extensive analysis, we demonstrate the zero-shot capabilities of CG3D, which enables zero-shot 3D classification, scene querying with natural language, and cross-modal retrieval. Furthermore, CG3D provides strong initial weights when training 3D networks for downstream tasks.

References

- [1] Mohamed Afham, Isuru Dissanayake, Dinithi Dissanayake, Amaya Dharmasiri, Kanchana Thilakarathna, and Ranga 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*, pages 9902–9912, 2022.
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022.
- [3] I. Armeni, A. Sax, A. R. Zamir, and S. Savarese. Joint 2D-3D-Semantic Data for Indoor Scene Understanding. *ArXiv e-prints*, Feb. 2017.
- [4] 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.
- [5] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [7] Yulong Chen, Yang Liu, Li Dong, Shuohang Wang, Chenguang Zhu, Michael Zeng, and Yue Zhang. Adaprompt: Adaptive model training for prompt-based nlp. *arXiv preprint arXiv:2202.04824*, 2022.
- [8] Jaemin Cho, Seunghyun Yoon, Ajinkya Kale, Franck Dernoncourt, Trung Bui, and Mohit Bansal. Fine-grained image captioning with clip reward. *arXiv preprint arXiv:2205.13115*, 2022.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [10] 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.
- [11] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. In *Advances In Neural Information Processing Systems*, 2022.
- [12] Shashank Goel, Hritik Bansal, Sumit Bhatia, Ryan A Rossi, Vishwa Vinay, and Aditya Grover. Cyclic: Contrastive language-image pretraining. *arXiv preprint arXiv:2205.14459*, 2022.
- [13] X Gu, T Lin, W Kuo, and Y Cui. Zero-shot detection via vision and language knowledge distillation. *corr abs/2104.13921* (2021).
- [14] 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.
- [15] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pages 709–727. Springer, 2022.
- [16] Linh Kästner, Vlad Catalin Frasinianu, and Jens Lambrecht. A 3d-deep-learning-based augmented reality calibration method for robotic environments using depth sensor data. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1135–1141. IEEE, 2020.
- [17] Konrad Koniariski and Andrzej Myśliński. Feature point cloud based registration in augmented reality. In *Advances in Systems Engineering: Proceedings of the 28th International Conference on Systems Engineering, ICSEng 2021, December 14-16, Wrocław, Poland 28*, pages 418–427. Springer, 2022.
- [18] Jie Lei, Linjie Li, Luwei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7331–7341, 2021.
- [19] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.
- [20] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021.
- [21] 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.
- [22] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019.
- [23] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *Advances in neural information processing systems*, 31, 2018.
- [24] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021.
- [25] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana

- Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. *arXiv preprint arXiv:2210.04150*, 2022.
- [26] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [28] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019.
- [29] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022.
- [30] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point cloud: A simple residual mlp framework. *arXiv preprint arXiv:2202.07123*, 2022.
- [31] 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.
- [32] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [33] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 4213–4220. IEEE, 2019.
- [34] Ron Mokady, Amir Hertz, and Amit H Bermano. Clip-cap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021.
- [35] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*, pages 529–544. Springer, 2022.
- [36] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [37] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. *arXiv preprint arXiv:2203.06604*, 2022.
- [38] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [43] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.
- [44] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- [45] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10529–10538, 2020.
- [46] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Point-rcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 770–779, 2019.
- [47] Kihyuk Sohn, Yuan Hao, José Lezama, Luisa Polania, Huiwen Chang, Han Zhang, Irfan Essa, and Lu Jiang. Visual prompt tuning for generative transfer learning. *arXiv preprint arXiv:2210.00990*, 2022.
- [48] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. Motionclip: Exposing human motion generation to clip space. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, pages 358–374. Springer, 2022.
- [49] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *International Conference on Computer Vision (ICCV)*, 2019.
- [50] Vibashan VS, Vikram Gupta, Poojan Oza, Vishwanath A. Sindagi, and Vishal M. Patel. Mega-cda: Memory guided attention for category-aware unsupervised domain adaptive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4516–4526, June 2021.
- [51] Vibashan VS, Domenick Poster, Suya You, Shuowen Hu, and Vishal M. Patel. Meta-uda: Unsupervised domain adaptive thermal object detection using meta-learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1412–1423, January 2022.
- [52] Vibashan VS, Jeya Maria Jose Valanarasu, and Vishal M Patel. Target and task specific source-free domain adaptive image segmentation. *arXiv preprint arXiv:2203.15792*, 2022.

- [53] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.
- [54] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019.
- [55] Zeyu Wang, Cuong Nguyen, Paul Asente, and Julie Dorsey. Pointshopar: Supporting environmental design prototyping using point cloud in augmented reality. 2023.
- [56] 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.
- [57] Jinheng Xie, Xianxu Hou, Kai Ye, and Linlin Shen. Clims: cross language image matching for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4483–4492, 2022.
- [58] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *European conference on computer vision*, pages 574–591. Springer, 2020.
- [59] Yinghui Xing, Qirui Wu, De Cheng, Shizhou Zhang, Guoqiang Liang, and Yanning Zhang. Class-aware visual prompt tuning for vision-language pre-trained model. *arXiv preprint arXiv:2208.08340*, 2022.
- [60] Chenfeng Xu, Shijia Yang, Bohan Zhai, Bichen Wu, Xiangyu Yue, Wei Zhan, Péter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Image2point: 3d point-cloud understanding with pretrained 2d convnets. *ArXiv*, abs/2106.04180, 2021.
- [61] 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.
- [62] Hengxu You, Fang Xu, and Eric Du. Robot-based real-time point cloud digital twin modeling in augmented reality. *Transforming Construction with Reality Capture Technologies*, 2022.
- [63] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen 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*, pages 19313–19322, 2022.
- [64] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- [65] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and language prompt learning. *arXiv preprint arXiv:2210.07225*, 2022.
- [66] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. *arXiv preprint arXiv:2112.02413*, 2021.
- [67] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [68] 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.
- [69] Chong Zhou, Chen Change Loy, and Bo Dai. Denseclip: Extract free dense labels from clip. *arXiv preprint arXiv:2112.01071*, 2021.
- [70] Peng Zhou, Rui Peng, Maggie Xu, Victor Wu, and David Navarro-Alarcon. Path planning with automatic seam extraction over point cloud models for robotic arc welding. *IEEE Robotics and Automation Letters*, 6(3):5002–5009, 2021.
- [71] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, pages 350–368. Springer, 2022.