

Received 4 March 2024, accepted 5 April 2024, date of publication 12 April 2024, date of current version 19 April 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3388155

RESEARCH ARTICLE

ExpPoint-MAE: Better Interpretability and Performance for Self-Supervised Point Cloud Transformers

IOANNIS ROMANELIS^{1,2}, VLASSIS FOTIS¹,
KONSTANTINOS MOUSTAKAS¹, (Senior Member, IEEE),
AND ADRIAN MUNTEANU², (Member, IEEE)

¹Department of Electrical and Computer Engineering, University of Patras, 265 04 Patras, Greece

²Department of Electronics and Informatics, Vrije Universiteit Brussel, 1050 Brussels, Belgium

Corresponding author: Ioannis Romanelis (iroman@ece.upatras.gr)

This work was supported by the European Union's Horizon 2020 Research and Innovation Program (DIDYMOS-XR: Digital DynaMic and responsible twinS for XR) under Grant 101092875.

ABSTRACT In this paper we delve into the properties of transformers, attained through self-supervision, in the point cloud domain. Specifically, we evaluate the effectiveness of Masked Autoencoding as a pretraining scheme, and explore Momentum Contrast as an alternative. In our study we investigate the impact of data quantity on the learned features, and uncover similarities in the transformer's behavior across domains. Through comprehensive visualizations, we observe that the transformer learns to attend to semantically meaningful regions, indicating that pretraining leads to a better understanding of the underlying geometry. Moreover, we examine the finetuning process and its effect on the learned representations. Based on that, we devise an unfreezing strategy which consistently outperforms our baseline without introducing any other modifications to the model or the training pipeline, and achieve state-of-the-art results in the classification task among transformer models.

INDEX TERMS Deep learning, explainability, point clouds, self-supervision, transformers.

I. INTRODUCTION

Deep learning models at large scale require adequately large labeled datasets to be able to learn. This is clearly a limitation for deep learning in general, since manual annotation is a very time-consuming and costly task. As a result, it was a great breakthrough when it was discovered that models can actually benefit from unlabeled data, by using them to design and solve pretext tasks, in which the sample itself is the label.

This concept appears to have features resembling a real-world analog, that is, the training happening inside a small child's/infant's brain. Supervision from experts (adults) is actually responsible for a small percentage of knowledge acquired throughout a human's lifespan, and it mostly tailors to specific material rather than the perception of the real world

The associate editor coordinating the review of this manuscript and approving it for publication was Walter Didimo¹.

at large. Simply by observing the behavior of other people and surrounding objects, they gain intuitive understanding of their environment. While we have fairly successfully replicated a form of supervised learning, self-supervision is an area that still lags behind, despite theoretically offering the most advantages.

Naturally, the data needs to be diverse and of sufficient size for this process to actually yield any benefits. Fortunately, by leveraging the ocean of unlabeled data that is freely available on the web, researchers have managed to significantly boost the performance of their models. Great examples of this concept can be found in both language [1] and in image [2], [3] domains, showing great promise for a variety of potential applications in other domains.

The most intuitive and (arguably) popular pretext task applied both in vision and language is the so-called "fill in the gaps". It refers to tasks where input samples are truncated

to occlude pieces of information or corrupted by applying various types of perturbations and noise. The deep learning model is given the damaged sample as input and is tasked to reconstruct/complete the original. During this process, the model learns features related to the sample's class. These features are incredibly useful and can be leveraged to improve the performance of the model in downstream supervised tasks, involving potentially much smaller datasets.

As this approach became more common, more sophisticated techniques of self-supervision were devised [4], [5], [6]. However, few studies have actually delved into what it is that the models actually learn [7], [8], [9] and even fewer have examined point clouds in particular [10], [11]. Point Cloud models differ significantly from image-based ones. As opposed to an image's canonical grid structure, uniform density and general data availability, the point clouds are highly irregular, highly inconsistent in terms of density, and are generally scarce compared to images. Additionally, contrary to an image's fixed pixel position, point clouds are unordered, often different in terms of cardinality, and require that their handling be invariant to permutations. Due to these challenges that arise in the domain shift from language/image to point clouds, we find that this particular area is lacking.

In this study we adopt the standard transformer as our baseline model. By "standard transformers" we refer to models that follow the architecture presented in ViT [12], as opposed to architectures that use transformer-like blocks. It is a highly versatile, strong architecture that has demonstrated incredible results in both language and vision, overtaking previous state-of-the-art models [13]. However, transformer-based works on point clouds do not reflect this success. In this work we explore various aspects of this architecture, in order to better understand its inner workings and find ways to improve its performance. We argue that our findings will be of further use in future works, since they are targeted at a widely used, general-purpose architecture rather than a specialized one.

To summarize, our contributions are twofold:

- We propose *strategic unfreezing*, a finetuning strategy that retains the properties of the backbone, learned through pretraining, while increasing the accuracy of our baseline both in ModelNet40 (+0.5%) and ScanObjectNN (+0.86%, +1.73%, +0.07%), and achieving state-of-the-art results among transformer models. (Sec. IV)
- We adjust explainability tools from the image and NLP domain to work with point clouds, with the aim of understanding the inner workings of the transformer and the effect of pretraining. Interestingly, we uncover that with more data, the transformer seems to learn the inductive bias of convolution, to attend locally. (Sec. V)

II. RELATED WORK

A. DEEP LEARNING ON POINT CLOUDS

The domain shift from 2D to 3D came with several extra challenges, due to the lack of grid structure and uneven density of point clouds. Early attempts involved

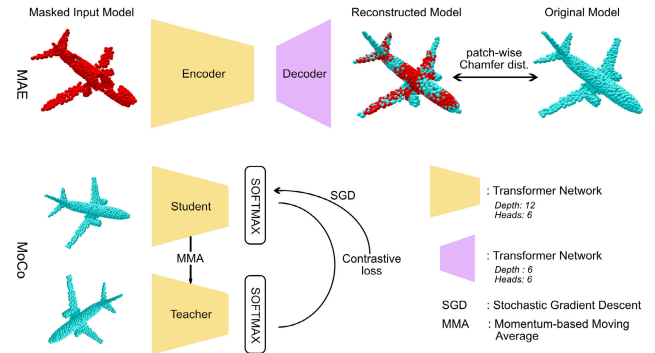


FIGURE 1. Graphical description of the two pretraining pipelines studied in this paper, namely Masked-AutoEncoding (MAE) [25] and Momentum Contrast (MoCo) [4]. In simple terms, MAE trains an autoencoder to reconstruct a shape with missing parts, whereas MoCo trains two networks (Student/Teacher) to generate approximately equal predictions for different augmentations of a data sample.

the application of image methods as-is, using multi-view images [14], [15] or voxels [16], [17]. Not long after that, specialized architectures emerged. PointNet [18] pioneered point-wise MLPs and pooling for extracting global features. PointNet++ [19] later followed a multi-scale approach by incorporating neighborhood information. Following this paradigm, other works create more complex kernels, taking advantage of geometric priors [20], [21], [22], [23], [24].

More recently, transformer variants have been successfully applied to the point cloud domain. Originally introduced in [26], transformers were quickly established as the go-to choice in NLP. Dosovitskiy et al. [12], later introduced the vision transformer, operating on image data. This was accomplished by tokenizing the input into patches and incorporating spatial information through a positional embedding. This approach also allowed data from different modalities to be used jointly [27], [28], making the transformer the most prominent architecture in that field. Point Transformer [29] is among the first works to apply a transformer-like architecture to point cloud data. They apply intra-neighborhood attention to create a patch feature vector and use pooling operations to downsample the point cloud.

In PCT [30], they extract point embeddings through a downsampling network and apply a series of single attention transformer layers to the remaining points. The feature maps of these layers are concatenated, and global features are extracted through mean and max pooling. In PVT [31], the authors create specialized attention modules for both points and voxels, utilizing the complementary information they contain to extract better features.

Point-BERT [32] tries to generalize the plain transformer, as used both in Image and NLP. They train a "tokenizer" autoencoding network, based on DGCNN [33], to map point neighborhoods into feature vectors. During pretraining, a transformer encoder-decoder architecture receives masked point clouds as input and tries to reconstruct the embeddings of the "tokenizer". Additionally, the classification token

is trained through a contrastive loss like in [4]. Point-MAE [25] and Point-Multiscale MAE (M2AE) [34] both utilize masked autoencoding to pretrain their transformer backbones, by reconstructing the actual points of the masked neighborhoods directly. Point-M2AE uses a pyramid-like backbone that gradually downsamples the input point cloud, obtaining multi-scale features. On the other hand, Point-MAE uses the same architecture as Point-BERT, that is, the plain transformer.

While hierarchical architectures such as M2AE are typically associated with slightly better performance, they require parameter tuning when transferred to other domains or scaled to larger datasets. Since the main scope of this work is explainability, the simplicity and cross-domain applicability of MAE are more valuable. It can be trivially scaled by adding or removing blocks and provides a fair ground for comparisons with image and NLP. Therefore, for our baseline model, we choose to adopt the transformer presented in Point-MAE [25].

B. SELF SUPERVISED LEARNING

The idea of making use of large amounts of unlabeled data to improve a learning model is not new, rather it has been around for decades [35], [36], [37], [38]. Its popularity did not surge until the mid-2010s [2], [3], [39] however, when the availability of unlabeled data at scale started to increase and the potency of hardware could finally keep up to the challenge. Of particular importance is masked autoencoding, pioneered by [1], which revolutionized language modeling. This paradigm was later followed by other works, such as [40], [41], and [42], that further scaled up in terms of parameter and dataset size, achieving state-of-the-art results in the few-shot and zero-shot settings. In [43] they experiment with pretraining by using synthetic data, artificially injected with desirable properties.

Another line of work that has been amassing popularity recently is contrastive learning [4], [5], [44], [45], [46], [47], [48] as a way to learn meaningful data representations. It is accomplished by generating positive and negative sample pairs and training the model to pull together or push them away, respectively. The pairs are generated through data augmentations; perturbations, crops, and transformations. A similar idea, dubbed contrastive clustering, is applied to groups of samples instead of pairs [6], [49]. Reference [6] in particular, eases up the computational load by introducing learnable cluster centers. In [50], the authors eliminate the need for negative pairs by training two networks to produce matching feature vectors for two different views of the same sample.

Despite the scarcity of point cloud data compared to image and language data, and the corresponding need for good self-supervision techniques in this domain, the topic has only recently begun attracting attention. Reference [51] proposes to split point clouds into parts and use these parts to parameterize gaussian mixture models. The model

is then trained by using a loss that resembles likelihood maximization. In [52] the authors perform patch-level contrastive learning. By rotating a query patch they form positive pairs, while negative pairs are formed by taking into account any other patch of the shape. A similar approach is followed in [53], finding positive and negative pairs by using an inductive model to perform pseudolabeling. In [54], a similar approach to masked autoencoding is followed, that is, completion of occluded point clouds acquired by taking custom viewpoints through virtual cameras. Reference [55] argues that performing pretext tasks on single object point clouds might have limited benefits in real world applications. They instead sample positive and negative pairs from complex scene scans, hoping to get a better estimate of the target distributions, and apply a contrastive loss to train their model.

C. EXPLAINABILITY IN DEEP LEARNING

An extremely important field in deep learning research is explainability. Deep models are generally seen as black boxes and any attempt at designing new or improving existing models is usually empirical. Nevertheless, there are several tools one can use in order to better understand a model's behavior. In [56] and [57] the authors present CKA, a similarity measure between feature representations of two arbitrary neural networks. Reference [7] utilizes this tool and provides valuable insights into the differences between how CNNs and transformers learn, while [8] tries to compare the representations of networks with different depths and feature map sizes. Other works [9], [11] utilize gradient-based methods in order to visualize the receptive fields or the relevancy of input patches towards the model's decision. We use a combination of all the above tools to attain a complete picture of our model and accompanying pretraining scheme.

III. BASELINE METHOD, TOOLS, AND DATASET

In this section, we briefly overview our baseline model and pretraining setup. After training and finetuning, we assess the quality of the learned features both qualitatively and quantitatively by using explainability tools and measuring the accuracy in the classification task. Further details regarding the model and the training process can be found in the appendix. We utilize pytorch lightning¹ for our codebase, in favor of clarity, transparency, and reproducibility. The code is publicly available on github.²

A. MASKED AUTOENCODERS

For our baseline model, we adopt the vision transformer trained with masked autoencoding, as presented by Pang et al. [25], which we restate here for the sake of completeness. The transformer is comprised of 12 blocks with 6 attention heads per block. The input point cloud is

¹<https://www.pytorchlightning.ai/>

²<https://github.com/VVRPanda/ExpPoint-MAE.git>




FIGURE 2. Comparison of different unfreezing points for the backbone of the transformer model. Unfreezing the model too early or too late can result in suboptimal results, as the network may either ‘forget’ the features learned through pretraining or fail to acquire task-specific knowledge, respectively.

split into N , possibly overlapping patches, 60% of which are masked out. The remaining patches are embedded into feature space via a small PointNet-like [18] network, and the patch centroids are used to generate positional encodings. The two are added to form the input to our transformer, which outputs a feature vector for each patch. At this stage, masked patches are assigned a mask token and a positional encoding and are concatenated with the rest before being fed into the decoder. The decoder reconstructs the masked patches and a reconstruction loss based on Chamfer distance is applied patch-wise.

B. EXPLAINABILITY TOOLS

For explainability, one of the main points of interest in this work, we utilize a variety of tools. The Centered Kernel Alignment criterion offers a versatile way of comparing feature representations between two models across a dataset, and its importance to our work cannot be overstated. It provides insights into the differences between models that have been pretrained & finetuned vs just finetuned, as well as models that have been pretrained using varying datasets and or strategies. We omit the math behind CKA and refer the interested reader to [56] and [57]. Attention visualization is another commonly used tool that helps to understand how the model correlates the point patches based on the extracted geometric features. Finally, we compute and visualize the receptive fields to better understand the information flow within the network.

C. DATASET

As opposed to their image counterparts, point cloud datasets are generally fewer and significantly smaller in size. Additionally, there are a lot of impactful differences between point cloud datasets, including sampling density, scanning device-specific artifacts, and shape variations in general (real-world objects, CAD models). As a result, a model

trained on a specific dataset may not generalize to other data distributions. Our first goal towards explainability is to see how the amount of training data might affect the network’s performance. To this end, we concatenate the commonly used ShapeNet [58] with the more obscure CC3D [59] (41k and 43k training samples, respectively). CC3D contains high resolution, single-object, fine-grained CAD shapes from arbitrary categories. To study the effects of the increased volume of data, we conduct pretraining experiments using only ShapeNet (S) and the concatenated dataset (C+S). We compare the two by finetuning for classification in ModelNet40 [60] and ScanObjectNN [61] and report the results in Tables 1, 2. In both cases, our proposed dataset helps achieve better accuracy. For the rest of the paper, we will be using C+S for pretraining, unless otherwise specified.

IV. STRATEGIC UNFREEZING

There are several ways to approach the finetuning process of a pretrained model. Most commonly, the pretrained backbone is frozen, that is, gradient propagation is allowed but no weight updates happen, and a smaller classification (or other) head is trained to map the learned features to logits. This method is based on the assumption that the backbone has already learned a robust representation of the data and is capable of separating the class clusters effectively. However, the learned data representation often differs significantly from the data used in the downstream tasks. A great example of this is pretraining on synthetic data and using the model in real-world applications.

Another approach is to perform task-specific adjustments to the model by training both the backbone and the added head end-to-end using a small learning rate. However this contains an important caveat, as the random initialization of the classification head’s weights can cause the weights of the backbone to be perturbed in an unintended direction, effectively destroying the learned features and representations.

In order to properly evaluate these approaches we perform comparative the following experiments.

- We test the first method by measuring the accuracies in ModelNet40 classification (Figure 2). We see that when keeping the backbone frozen and only relying on the classification head, the accuracy is vastly inferior compared to the backbone being unfrozen at various intervals. This is because no task-specific knowledge has been incorporated into the backbone.
- In Figure 3(a), we test the second method. We compare the pretrained backbone that has been finetuned without freezing and the pretrained backbone without finetuning. It is evident that there is hardly any similarity between them, except for the very early layers. This arguably defeats the purpose of pretraining, since very few of the valuable properties are retained in the final network. This deficiency is also reflected in the final accuracy score, as seen in Figure 2.
- Finally, in Figure 2 we note the accuracy of the model that has not been pretrained. It exhibits the lowest




FIGURE 3. CKA comparison of the pretrained backbone (y-axis) with versions that have been finetuned for 300 epochs, unfreezing the backbone on various epochs (x-axis). The first and second blocks indicate the positional and feature embedding extractors, while the rest of the blocks correspond to the outputs of the attention layers. High values indicate high similarity between feature representations. The later the network is unfrozen, the higher the similarity with the pretrained backbone, retaining the properties learned through pretraining. In the case of (a), as done in [25], the final network has very little similarity with the pretrained backbone, thereby nullifying the effectiveness of pretraining.

TABLE 1. Comparison of different transformer based models and pretraining schemes with and without using voting on the ModelNet40 test set.

MODEL	W/O VOTING	VOTING
Point-BERT [32]	92.7	93.2
MaskPoint [62]	-	93.8
Point-MAE [25]	93.2	93.8
PCT-2L [30]	93.2	-
PCT-3L [30]	93.4	-
Point-M2AE [63]	93.4	94.0
Ours (S)	93.6	-
Ours (C+S)	93.7	94.2

Blue and Cyan denote 1st and 2nd highest.

accuracy among all models that have been pretrained, proving the effectiveness of pretraining.

Before diving into what the model actually learns (section VI), we first propose a two-stage training scheme that retains the pretrained network’s characteristics and obtains superior performance. First, we train the classification head for the majority of epochs, ensuring that it has learned to separate the clusters formed via the backbone. Finally, we jointly train both of them for the remainder of the epochs. In order to figure out the appropriate step to unfreeze the backbone, we perform several experiments and compare the feature representations with the pretrained network each time (Figure 3). Additionally, we measure the best and the convergence accuracy for each of these models in the classification task on ModelNet40 (Figure 2). It is evident that the earlier the backbone is unfrozen, the more the pretrained backbone’s features are distorted, and the final accuracy is compromised.

In order to identify the proper unfreezing epoch, one must consider the data and the task. For ModelNet40 classification, as a subset of ShapeNet, similar samples were seen by the model during pretraining. As a result, unfreezing during the late stages is appropriate. On the other hand, when finetuning on ScanObjectNN, that consists of real scanned objects, it is wiser to unfreeze earlier so that the backbone can become accustomed to the new data distribution. Finally, when finetuning on the hardest variant of the ScanObjectNN which includes other modifications, such as rotations, we found that the most effective way is to perform an extra domain

TABLE 2. Classification results on ScanObjectNN. The HARD variant refers to PB-T50-RS that is the most challenging one.

MODEL	OBJ-BG	OBJ-ONLY	HARD
Point-BERT	87.43	88.12	83.07
MaskPoint	89.3	89.7	84.6
Point-MAE	90.02	88.29	85.18
Point-M2AE	91.22	88.81	86.43
Ours (unfr 200)	90.88	89.33	84.98
Ours (unfr 250)	87.44	90.02	83.80
Ours (SoNN ¹)	90.36	89.67	85.25

Blue and Cyan denote 1st and 2nd highest.

¹: SoNN indicates that the model has been pretrained for a few epochs in the ScanObjectNN-hard dataset.

adaptation pretraining step, before finetuning. The results of this step can be seen in Table 2 in the SoNN entry. More details are presented in the Appendix.

Designing and training a model is a difficult and iterative process, that is based heavily on parameter tuning. We believe that with our method researchers will be able to make more informed decisions that will allow them to utilize their models to their maximum capacity.

V. EXPLAINABILITY STUDY

A. ATTENTION VISUALIZATION

In this section, we will use the finetuned backbone, pretrained with MAE, and finetuned with the strategy outlined in the previous section. This is because it retains the useful properties acquired through pretraining, but also has a trained classification token, as opposed to [25], which uses max-average pooling. By visualizing the attention matrices at each block we can obtain meaningful information regarding how the model correlates various parts of the shape. This information can help identify patterns that indicate whether the pretraining procedure was successful, such as semantic correlation, symmetry, and locality.

We observe that in all input shapes, the first attention block is always global, while the rest gradually shift from attending to specific regions to attending globally. Sharp geometric features in particular attract high attention scores in most layers (Figure 4). This behavior is quantitatively verified through attention distances in the following section.




FIGURE 4. Attention Visualization of the classification token for each block (1-12), averaged across heads (brighter = higher score). Although the classification token’s attention score towards itself cannot be visualized, it has the highest value in all cases. This score is included in the normalization process, so that the relative scale between them is visible.



FIGURE 5. Attention Visualization of the classification token for each block (1-12), averaged across heads for a finetuned model without pretraining. It is evident the locations it attends to do not follow any recognizable pattern. In each layer, the attention is focused on a specific location, which is often the same between layers (e.g. head in layers 0,5,7,10).

It is important to note that finetuned models that have not been pretrained **do not** share this trait. As can be seen in figure 5, although the classification token attends to locations with valid semantic meaning, the behavior is erratic, repetitive and there is no clear transition from specific areas to the whole shape.

B. ATTENTION DISTANCE

Following the paradigm of [7], we study the sorted attention distance. The mean Attention Distance (mAD) of layer l for a head h is calculated as follows:

$$mAD_{l,h} = \frac{1}{N^2} \sum_{ij} \mathbf{A}_l \odot \mathbf{D}$$

where $\mathbf{A}_l^{N \times N}$ is the attention matrix of layer l and $\mathbf{D}^{N \times N}$ is the distance matrix, where each entry $d_{i,j}$ is the Euclidean distance between centers i and j and the \odot symbol denotes the Hadamard product. The product is averaged across all entries to obtain a mean attention-distance value, which is also averaged across the validation set of ModelNet40.

Based on the graphs in Figure 6, it can be inferred that with the increase in the amount of training data, the earlier layers tend to focus more locally with certain attention heads, while others still maintain a global focus. This observation suggests that the transformer network may begin to learn the inductive bias of the convolution, that is, attending to local features, while still retaining the capability to aggregate global information, thereby resulting in more comprehensive representations. This mitigates the general struggle of self-attention to model local relationships [64], without requiring any modifications to the architecture.

Interestingly, our findings are consistent with those of [7], [65], and [66], indicating that transformers may exhibit akin

behaviors across domains. A noticeable difference is that in the case of ShapeNet (Figure 6b) the first two layers seem attend globally with all of their attention heads, whereas when CC3D is added to the pretraining data, this phenomenon is only observed in the first layer (Figure 6c). We assume that the network requires a rough understanding of the entire shape before opting to extract local features. This demeanor should be studied further if point cloud datasets at large scale become available in the future.

C. EFFECTIVE RECEPTIVE FIELDS

We have established that by using our unfreezing strategy, the final network shares many similarities with the pretrained network. In order to understand their differences (trained for pretext versus downstream tasks), we visualize the effective receptive fields [7]. We accomplish this by selecting a patch and propagating the gradients backwards from its feature vector to its embedding, including both feature and positional information. We then visualize the norm of said gradients and present the results in figure 7.

We notice that the pretrained model, trained explicitly for reconstruction, looks strictly at neighboring as well as semantically similar patches, since they contain the most relevant information. These areas have the highest impact on the finetuned model as well, but its receptive field is noticeably wider. We believe this is because the classification task requires global features.

VI. CONTRASTIVE LEARNING

Having delved into the inner workings of masked autoencoding and the properties associated with it as a pretraining scheme, we need to evaluate its effectiveness against other pretraining methods in the literature. We choose what we




FIGURE 6. Attention Distance of the attention heads (distances are sorted in ascending order for clarity). The distances are averaged across the entire validation set of ModelNet40. We see that as the amount of data increases, early layers of the network begins to attend locally, while higher layers incorporate global information, exhibiting similar behaviour to ViT in the image domain [12] Notably, a key difference is observed in the very first layer, where all heads consistently attend globally.




FIGURE 7. Effective receptive fields of the patch of interest (colored in red). We compute the gradient of the feature representations generated by the last layer with respect to the input patches. Each point is colored based on the magnitude of the gradient of its corresponding patch. In the right variant, that is finetuned for classification, the model retrieves information from the entire shape, as it tries to incorporate class information. Conversely, the left variant, that is the pretrained backbone, focuses mostly to local and symmetric parts, aligning with the task it was trained for, that is reconstruction.



FIGURE 8. Attention scores of the classification token, averaged across all heads, pretrained using the MoCo pipeline. We visualize a subset of the network layers, [1, 3, 7, 8] left to right. While in the early layers the CLS token attends to meaningful parts of the shape, exhibiting a tendency to attend at symmetric parts, from layer 8 and onwards it attends uniformly to the whole shape.




FIGURE 9. Comparison of Attention Distances achieved through Masked AutoEncoding (MAE) [25] and Momentum Contrast (MoCo) [4] pretraining schemes. We observe that the setup of MoCo is insufficient to enable the transformer to learn to attend locally. Layers 10 and 11 in MoCo have the exact same attention distance.

believe to be the next most popular one, *contrastive learning*. Impressed by its performance in the image domain, we follow the approach of *Momentum Contrast (MoCo)* [4], which we extend to work with 3D data.

In MoCo, two copies of the network are trained simultaneously, one through backpropagation and the other using a momentum update rule. Positive pairs are created by feeding crops of the same shape with different augmentations to these networks, while previous activations of the momentum network act as negative pairs. We follow the same strategy and create our crops by using a mix of gaussian noise, anisotropic scaling, rotations, and random point dropouts.

By finetuning the models pretrained with MoCo, we observe that, although it improves the original model without pretraining, there is a significant gap in accuracy compared to its MAE counterpart (Table 3). This, however, does not rule out the possibility that other variants of contrastive learning might be more effective. Instead, it suggests that a well-designed pretraining scheme will most

TABLE 3. Network accuracy achieved with different pretraining schemes.

Model	No Pre.	MAE	MoCo	MAE & MoCo
Accuracy	91.1	93.7	92.2	91.8

likely offer a significant performance boost to any baseline model.

As opposed to masked autoencoding, an interesting property of contrastive learning is that it allows the classification token to be trained during the pretraining stage. In fact, by examining the attention scores, we notice that the classification token attends to characteristic areas of the input




FIGURE 10. KCA comparisons between the two pretrained backbones (MAE and MoCo pipelines) and themselves. The first two entries correspond to the positional and feature embedding respectively while the rest to the outputs of the attention blocks. We observe little resemblance between the representations learned through the different pretraining pipelines. Interestingly, the final feature representations obtained through the MAE pipeline appear to rely more on the features learned within the layers of the network, showing little similarity with the positional embeddings. Conversely, the MoCo pipeline demonstrates a higher reliance on the positional embedding.

shape, and exhibits symmetry, by attending to shape parts with similar semantics (Figure 8). We also observe that from layers 8 and onward, the CLS token attends uniformly to the whole shape. This is verified by Figure 9b where layers 10 and 11 have the exact same attention distance across all heads. Finally, we see that through MoCo pretraining the backbone does *not* learn to attend to neighboring patches, indicating that the training scheme might require more data or alterations to converge. For the sake of completeness, we also tested a MoCo and MAE hybrid (details in the Appendix).

Furthermore, useful insights can be extracted by comparing the two pretraining schemes against each other, as well as themselves, to get an idea of how the information is distributed throughout the layers. The results can be seen in Figure 10.

We observe that middle layers in MAE seem to have similar feature representations, while there is some similarity between input and output layers. However, the same does not seem to apply to MoCo, where each layer has common representations only with a few of its preceding and succeeding layers, creating a chain-like representation. Astoundingly, the layers seem to gradually lose their similarity with the early feature embedding and rely more and more on the positional embedding, whereas in MAE, the feature embedding is the most dominant in all layers. Naturally, MAE and Moco are based on fundamentally different principles, so they bear very little resemblance and only in the middle layers.

VII. ABLATION STUDIES

In the ablation study section we focus on two separate topics that are the masking ratio and the intermediate representations of the transformer layers.

Firstly, we focus of the reconstruction abilities of the backbone, trained through the MAE pipeline. Point-MAE has already studied the effect of masking in the final accuracy, concluding that a 60% masking is the optimal choice. In our experiments we want to test the reconstruction ability of the network when masking the input with a higher ratio. As depicted in Figure 11, even in the extreme case when 90%




FIGURE 11. Reconstruction results using different masking ratios. The network is pretrained, using MAE, with a 60% masking ratio. We test the reconstruction ability of the system when we provide as input a point cloud with higher masking ratios. Input and reconstructed patches are shown in red and cyan, respectively.

TABLE 4. Classification results on ModelNet40 when using the CLS token from previous layers as global descriptor of the 3D shapes. The network is trained with our unfreezing strategy, that is training only the cls head for the first 250 epochs.

Model	Point-BERT	Point-MAE	Ours(full depth)		
Acc.	92.70	93.19	93.70		
Without Layers	-5	-4	-3	-2	-1
Acc.	93.23	93.56	93.48	93.40	93.52

of the model is masked out, the network can provide a decent reconstruction. We believe that this is because positional embeddings of the masked patches are given to the decoder, providing useful hints regarding the position of individual parts.

The classification token in the transformer model progressively accumulates information in order to get a high-level understanding of the input shape. In our second experiment we try to evaluate the representation attained through intermediate layers. In Figure 12 we use the same pretrained backbone, but use only a subset of its transformer blocks. Numeric results are presented in Table 4. When removing up to 5 layers, the model manages to retrain an accuracy above the baseline. Interestingly, that’s the same layer




FIGURE 12. Accuracy graph of truncated copies of the same model. Each entry specifies the layer from which the feature vector is derived (all succeeding layers are discarded for that run). We observe that the model maintains higher accuracy than the baseline, even with a fraction of the original model's parameters.

in which the transition from location-specific to uniform attention becomes clearer. This indicates that the class cluster separation is possible even in earlier layers, and any additional information simply contributes to refining these clusters.

VIII. DISCUSSION

Using our proposed unfreezing strategy, we manage to outperform our baseline model, as well as models that use transformers in a similar way. We do not use voting unless explicitly mentioned for fair comparisons, nor do we include random transforms in the validation set. The method we suggest combined with our learning rate scheduling policy yield stable training with very close convergence points and weaker fluctuations.

Given sufficient training time, task-specific knowledge must be incorporated into the backbone in order to improve performance, but only when the head network has been trained enough for the weights to move in the appropriate direction. The specific intervals can vary depending on the task and dataset. For instance, data with background obtained from scans (in ScanObjectNN) require more train time for the backbone, so the optimal unfreezing point is around 200 epochs instead of 250. We also find that when the data distribution is fairly different from the pretraining data, it is more beneficial to pretrain our network for a few epochs in the new distribution instead of unfreezing the backbone earlier (Table 2).

In our explainability experiments, we observe that point neighborhoods with sharp geometric features are generally more likely to be attended to by the classification token. This is an indicator of good pretraining, as sharp features are more relevant descriptors of the sample's class. We also notice that the transition from local to global seems characteristic of transformers. This property is experimentally backed in the image domain [7], and our own experiments confirm that this

holds true in point clouds as well. An additional surprising finding is the contrast between the behaviors of two finetuned models, one that has been pretrained and one that has not. In the first case the human eye can clearly capture a transition from local to global, whereas in the second case the attention scores are mostly random, containing no identifiable pattern.

Finally, we remark that doubling the volume of data causes more attention heads to attend locally and introduces accuracy boosts in the classification task, albeit small. In [7] they observe similar behavior but with even greater impact, which is likely attributed to the significantly larger scale of both the datasets and models. Nevertheless, in both cases it suggests that local information aggregation is a property linked with good performance.

IX. FUTURE WORK

Our analysis in section VI demonstrates that contrastive learning exhibits useful properties, namely, symmetry and the ability to train a classification token in the pretraining stage. Consequently, we believe that contrastive learning can potentially outperform MAE in several tasks, given enough training time and data. In fact, we theorize that due to their distinct characteristics, some pretraining schemes may offer greater advantages for a specific downstream task than others. In future work, we plan on experimenting on a framework for determining which pretraining scheme to use, in order to utilize our models to their fullest potential.

APPENDIX A

MOCO & MAE HYBRID PRETRAINING

We devise a pretraing sceme that combines both core concepts from MoCo and MAE. The main idea is that the reconstruction loss from MAE will help the backbone converge faster, being a more straightforward task, and the contrastive loss will help in training the CLS token and slightly contribute to the backbone features as well. The scheme is realized through a weighted sum of both losses, appropriately scaled so that they are in the same order of magnitude.

Specifically, the standard reconstruction loss presented in [33] is applied (L_{rec}) and a contrastive loss, L_{con} , as in [20] is added, using a queue size of 4096 encoded samples. The final loss function is realized as:

$$L = L_{rec} + w_c \cdot L_{con}$$

where w_c was set to 10^{-2} to bring the two losses to the same order of magnitude. We also experimented with a smaller weight (10^{-3}), as the MoCo loss is responsible mainly for the CLS token. Both approaches, however, produced similar results.

Although the hybrid model's accuracy reaches an unimpressive 91.8% score, we believe that with further research it is possible to create a model that makes the best of both worlds (masked autoencoding and contrastive learning).

A. STRATEGIC UNFREEZING

As briefly discussed in the main paper (due to space constraints), the finetuning data from ModelNet40 have already been seen during the pretraining stage and therefore the distribution of said data has been captured by the backbone. Therefore, it is sufficient to only train the backbone for the last 50 epochs, where the learning rate is small enough so that no overblown weight updates happen.

The ScanObjectNN is probably an even better indicator of this strategy. We observe that in table 2, when training on objects without background the same unfreezing point yields the best results. On the other hand, when background is included, unfreezing at 200 epochs is optimal and obtains even better accuracy. While the higher accuracy can easily be attributed to the extra information contained in the background itself, it is important to notice that the background changes the data distribution significantly. This change, coupled with the fact that the data has been obtained through scanning, leads us to think that 50 epochs is no longer enough for the backbone to adjust to the new distribution.

This argument is further verified by finetuning on the hard variant of ScanObjectNN. Sure enough, 100 epochs of training with the unfrozen backbone makes a staggering +1.18% difference in accuracy compared to 50 epochs, however it is still far from the optimal, which is to *pretrain* the network for a few epochs on this dataset. In our experiments, 40 epochs of pretraining are adequate to reach an accuracy of 85.25%, that surpasses our baseline Point-MAE.

Generally speaking, there is no explicit way to determine the exact epoch, at which the unfreezing yields optimal results. However, our proposed “strategic unfreezing” method is a practical way of iteratively improving the model at hand, by providing concrete proof regarding the similarity with the original pretrained version of the model. Consequently, it allows the researcher to make informed design choices and gain an overall deeper understanding.

B. OPTIMIZATION AND LEARNING RATE POLICY

For both our optimizer and learning rate we follow one of [33]’s suggestions. For our finetuning we use the *AdamW* optimizer with a peak learning rate of $\lambda = 5 \cdot 10^{-4}$ and a weight decay of $W_d = 0.05$. We utilize the *cosine annealing learning rate scheduler*, updating during every step for a total training of 300 epochs. A *linear warmup* is also performed, with an initial learning rate $\lambda_i = 10^{-6}$ for a duration of 10 epochs in total. The exact same setup is also used for pretraining, with the exception of peak learning rate, being $\lambda = 10^{-3}$. The batch size is 128 for pretraining and 32 for finetuning.

C. DATA PREPROCESSING

All data samples are normalized to be inside the unit sphere. For data augmentations we apply random sampling of the point clouds and anisotropic scale, following the paradigm of our baseline.

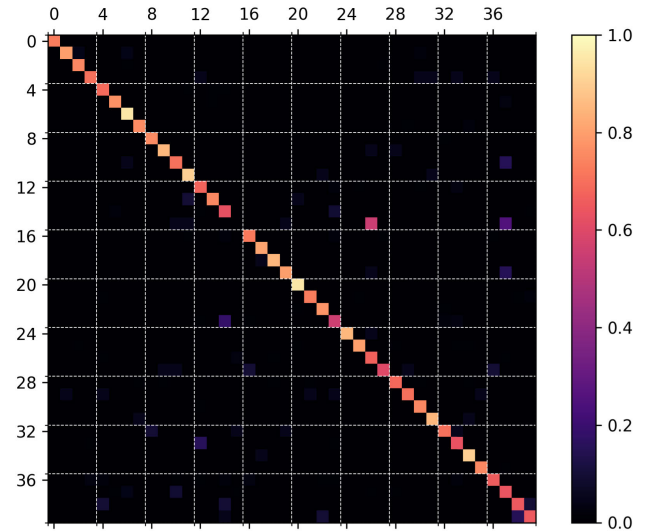


FIGURE 13. Confusion matrix visualization for the validation set of ModelNet40. Each row has been normalized through division by the number of samples of that class.




FIGURE 14. Visualization of samples from problematic classes of ModelNet40. First, second and third rows correspond to “flower pot”, “plant” and “vase” respectively.

D. HARDWARE

All models were trained on a single *NVIDIA RTX 4090* GPU. It should be noted that the memory requirements did not exceed 16GB. The training times were (approximately) 16 hours for pretraining on the combined dataset. For finetuning on our best model (unfreeze 250) the training time was 2 hours on ModelNet40 and 1 hour on ScanObjectNN. Note that when unfreezing at an earlier stage, the total time

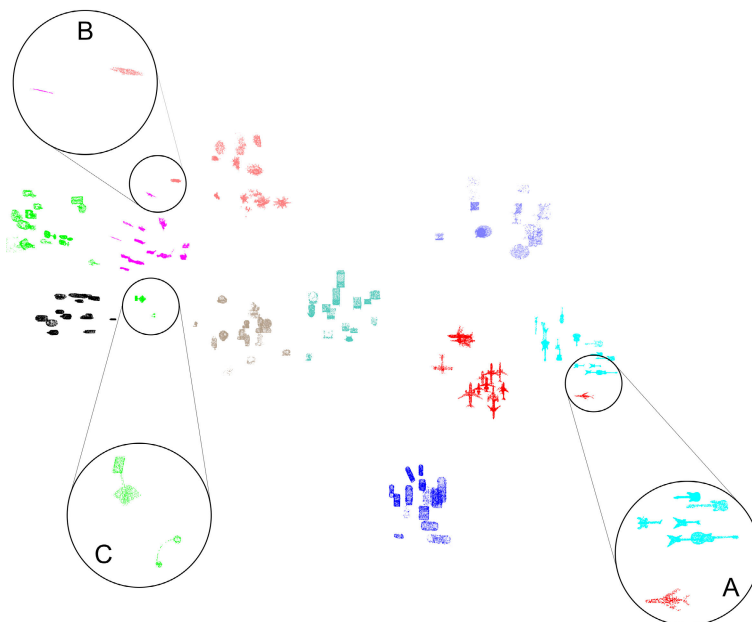


FIGURE 15. Clustering performed by the model in the feature space for 10 arbitrary classes. The model has been pretrained using the MAE pipeline. T-SNE is used to project the models two the 2-dimensional space. (A) A sample from the *airplane* class is placed near the *guitar* cluster. An “excusable” choice as both share a lengthy shape and sharp protruding parts. (B) Two shape from the *person* and *plant* classes are placed further from their clusters, and near each other, as they have a near non-existent depth. (C) Two uniquely shaped lamps have been assigned to a cluster of their own.

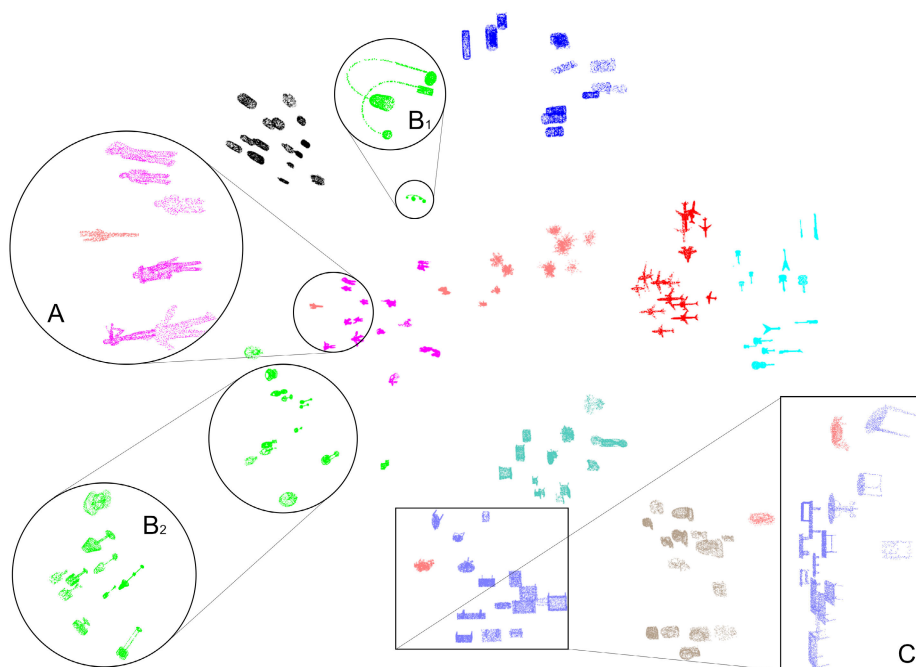


FIGURE 16. Clustering performed by the model in the feature space for 10 arbitrary classes. The model has been pretrained using the MoCo pipeline. T-SNE is used to project the models two the 2-dimensional space. (A) A tall cactus is clustered together with the *people* class. (B1), (B2), Two lamps with distinct characteristics have been assigned to a separate cluster. (C) A plant with hanging leaves is placed along with the *table* cluster, most likely because the leaves share a similar structure with table legs.

is lengthened due to the weight updates happening inside the backbone. It takes the model approximately 36 ms to process a batch with 16 samples.

APPENDIX B CLASS CLUSTERS

In order to get a better idea of the causes behind misclassified objects we present a confusion matrix (figure 13) generated from samples of ModelNet40 validation set. The most striking entry is the 15th category, “flower pot”, being mistaken for classes 26, “plant”, and 37, “vase”. These categories contain an unreasonably large amount of overlap, e.g. pots have plants, plants have pots and a lot of pots have a vase-like shape (figure 14). Therefore, it is hardly a surprising result that these objects are being misclassified, and this particular error is not a good descriptor of the model’s performance.

In figures 15 and 16 (MAE and MoCo respectively) we visualize the feature vectors of 15 shapes from 10 arbitrary classes randomly sampled from the validation set of ModelNet40 and projected onto the 2d space by using T-SNE. Naturally, the test accuracy is the sole indicator of how well the clustering is performed. However, by using this technique we are hoping to understand whether the clustering/classification operates within the realms of human logic. For this purpose, we are mostly interested in the misclustered samples.

All in all, these mistakes seem plausible from a human perspective. In particular, the model seems to be “fooled” by objects that have similar height/width or are composed by roughly equivalent parts. This is hardly a fault however, as it only decreases the score on the ModelNet40 benchmark, whilst achieving greater generalization ability.

ACKNOWLEDGMENT

(Ioannis Romanelis and Vlassis Fotis contributed equally to this work.)

REFERENCES

- [1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, vol. 1. Minneapolis, MI, USA: Association for Computational Linguistics, Jun. 2019, pp. 4171–4186.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” *Commun. ACM*, vol. 60, no. 6, pp. 84–90, May 2017.
- [3] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [4] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, “Momentum contrast for unsupervised visual representation learning,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 9726–9735.
- [5] T. Chen, S. Kornblith, M. Norouzi, and G. E. Hinton, “A simple framework for contrastive learning of visual representations,” in *Proc. 37th Int. Conf. Mach. Learn.*, vol. 119, 2020, pp. 1597–1607.
- [6] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, “Unsupervised learning of visual features by contrasting cluster assignments,” in *Proc. NIPS*, Dec. 2020, pp. 9912–9924.
- [7] M. Raghu, T. Unterthiner, S. Kornblith, C. Zhang, and A. Dosovitskiy, “Do vision transformers see like convolutional neural networks?” in *Proc. Conf. Neural Inf. Process. Syst. (NIPS)*, vol. 34, 2021, pp. 12116–12128.
- [8] T. Nguyen, M. Raghu, and S. Kornblith, “Do wide and deep networks learn the same things? Uncovering how neural network representations vary with width and depth,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2021, pp. 1–25.
- [9] H. Chefer, S. Gur, and L. Wolf, “Transformer interpretability beyond attention visualization,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 782–791.
- [10] B. Zhang, S. Huang, W. Shen, and Z. Wei, “Explaining the PointNet: What has been learned inside the PointNet?” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2019, pp. 71–74.
- [11] J. Tayyub, M. Sarmad, and N. Schönborn, “Explaining deep neural networks for point clouds using gradient-based visualisations,” in *Proc. Asian Conf. Comput. Vis. (ACCV)*, 2022, pp. 2123–2138.
- [12] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16×16 words: Transformers for image recognition at scale,” in *Proc. ICLR*, 2021.
- [13] X. Chen, C.-J. Hsieh, and B. Gong, “When vision transformers outperform ResNets without pre-training or strong data augmentations,” 2021, *arXiv:2106.01548*.
- [14] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller, “Multi-view convolutional neural networks for 3D shape recognition,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 945–953.
- [15] C. R. Qi, H. Su, M. Nießner, A. Dai, M. Yan, and L. J. Guibas, “Volumetric and multi-view CNNs for object classification on 3D data,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 5648–5656.
- [16] H. You, Y. Feng, R. Ji, and Y. Gao, “PVNet: A joint convolutional network of point cloud and multi-view for 3D shape recognition,” in *Proc. 26th ACM Int. Conf. Multimedia*, Oct. 2018, pp. 1310–1318.
- [17] D. Maturana and S. Scherer, “VoxNet: A 3D convolutional neural network for real-time object recognition,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2015, pp. 922–928.
- [18] R. Q. Charles, H. Su, M. Kaichun, and L. J. Guibas, “PointNet: Deep learning on point sets for 3D classification and segmentation,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 77–85.
- [19] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “PointNet++: Deep hierarchical feature learning on point sets in a metric space,” in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Red Hook, NY, USA: Curran Associates, 2017, pp. 5105–5114.
- [20] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. Guibas, “KPConv: Flexible and deformable convolution for point clouds,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6410–6419.
- [21] Y. Liu, B. Fan, S. Xiang, and C. Pan, “Relation-shape convolutional neural network for point cloud analysis,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 8887–8896.
- [22] X. Ma, C. Qin, H. You, H. Ran, and Y. Fu, “Rethinking network design and local geometry in point cloud: A simple residual MLP framework,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2022, pp. 1–8.
- [23] T. Xiang, C. Zhang, Y. Song, J. Yu, and W. Cai, “Walk in the cloud: Learning curves for point clouds shape analysis,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 895–904.
- [24] Y. Xu, T. Fan, M. Xu, L. Zeng, and Y. Qiao, “SpiderCNN: Deep learning on point sets with parameterized convolutional filters,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, vol. 11212, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds. Cham, Switzerland: Springer, 2018, pp. 90–105.
- [25] Y. Pang, W. Wang, F. E. H. Tay, W. Liu, Y. Tian, and L. Yuan, “Masked autoencoders for point cloud self-supervised learning,” in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 604–621.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proc. Conf. Neural Inf. Process. Syst. (NIPS)*, vol. 30, 2017.
- [27] A. Shin, M. Ishii, and T. Narihira, “Perspectives and prospects on transformer architecture for cross-modal tasks with language and vision,” *Int. J. Comput. Vis.*, vol. 130, no. 2, pp. 435–454, Feb. 2022.
- [28] X. Wu and T. Li, “Sentimental visual captioning using multimodal transformer,” *Int. J. Comput. Vis.*, vol. 131, no. 4, pp. 1073–1090, Apr. 2023.

- [29] H. Zhao, L. Jiang, J. Jia, P. Torr, and V. Koltun, "Point transformer," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 16239–16248.
- [30] M.-H. Guo, J.-X. Cai, Z.-N. Liu, T.-J. Mu, R. R. Martin, and S.-M. Hu, "PCT: Point cloud transformer," *Comput. Vis. Media*, vol. 7, no. 2, pp. 187–199, Jun. 2021.
- [31] C. Zhang, H. Wan, X. Shen, and Z. Wu, "PVT: Point-voxel transformer for point cloud learning," *Int. J. Intell. Syst.*, vol. 37, no. 12, pp. 11985–12008, Dec. 2022.
- [32] X. Yu, L. Tang, Y. Rao, T. Huang, J. Zhou, and J. Lu, "Point-BERT: Pre-training 3D point cloud transformers with masked point modeling," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 19291–19300.
- [33] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph CNN for learning on point clouds," *ACM Trans. Graph.*, vol. 38, no. 5, pp. 1–12, Oct. 2019.
- [34] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 15979–15988.
- [35] Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and L. Jackel, "Handwritten digit recognition with a back-propagation network," in *Proc. Conf. Neural Inf. Process. Syst. (NIPS)*, 1989.
- [36] A. Ng and M. Jordan, "On discriminative vs. generative classifiers: A comparison of logistic regression and naive Bayes," in *Proc. Conf. Neural Inf. Process. Syst. (NIPS)*, 2002.
- [37] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng, "Self-taught learning: Transfer learning from unlabeled data," in *Proc. 24th Int. Conf. Mach. Learn.*, Jun. 2007, pp. 759–766.
- [38] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proc. 25th Int. Conf. Mach. Learn. - ICML*, 2008, pp. 1096–1103.
- [39] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, May 2015.
- [40] T. B. Brown, "Language models are few-shot learners," in *Proc. NIPS*, 2020, pp. 1877–1901.
- [41] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," GPT-2, OpenAI, San Francisco, CA, USA, Tech. Rep., 2018.
- [42] X. Zhai, X. Wang, B. Mustafa, A. Steiner, D. Keysers, A. Kolesnikov, and L. Beyer, "LiT: Zero-shot transfer with locked-image text tuning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 18102–18112.
- [43] H. Kataoka, K. Okayasu, A. Matsumoto, E. Yamagata, R. Yamada, N. Inoue, A. Nakamura, and Y. Satoh, "Pre-training without natural images," *Int. J. Comput. Vis.*, vol. 130, no. 4, pp. 990–1007, Apr. 2022.
- [44] X. Chen, S. Xie, and K. He, "An empirical study of training self-supervised vision transformers," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 9620–9629.
- [45] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," 2020, *arXiv:2003.04297*.
- [46] M. Caron, H. Touvron, I. Misra, H. Jegou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 9630–9640.
- [47] L. Jiang, S. Shi, Z. Tian, X. Lai, S. Liu, C.-W. Fu, and J. Jia, "Guided point contrastive learning for semi-supervised point cloud semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 6403–6412.
- [48] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," 2018, *arXiv:1807.03748*.
- [49] M. Caron, P. Bojanowski, A. Joulin, and M. Douze, "Deep clustering for unsupervised learning of visual features," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018.
- [50] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. G. Azar, B. Piot, k. kavukcuoglu, R. Munos, and M. Valko, "Bootstrap your own latent—A new approach to self-supervised learning," in *Proc. Conf. Neural Inf. Process. Syst. (NIPS)*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33, 2020, pp. 21271–21284.
- [51] B. Eckart, W. Yuan, C. Liu, and J. Kautz, "Self-supervised learning on 3D point clouds by learning discrete generative models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 8244–8253.
- [52] B. Du, X. Gao, W. Hu, and X. Li, "Self-contrastive learning with hard negative sampling for self-supervised point cloud learning," in *Proc. 29th ACM Int. Conf. Multimedia*, Oct. 2021, pp. 3133–3142.
- [53] A. Cheraghian, S. Rahman, T. F. Chowdhury, D. Campbell, and L. Petersson, "Zero-shot learning on 3D point cloud objects and beyond," *Int. J. Comput. Vis.*, vol. 130, no. 10, pp. 2364–2384, Oct. 2022.
- [54] H. Wang, Q. Liu, X. Yue, J. Lasenby, and M. J. Kusner, "Unsupervised point cloud pre-training via occlusion completion," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 9762–9772.
- [55] S. Xie, J. Gu, D. Guo, C. R. Qi, L. Guibas, and O. Litany, "PointContrast: Unsupervised pre-training for 3D point cloud understanding," in *Proc. Eur. Conf. Comput. Vis. Cham, Switzerland: Springer*, 2020, pp. 574–591.
- [56] K. Cortes, M. Mohri, and A. Rostamizadeh, "Algorithms for learning kernels based on centered alignment," *J. Mach. Learn. Res.*, vol. 13, no. 1, pp. 795–828, Mar. 2012.
- [57] S. Kornblith, M. Norouzi, H. Lee, and G. Hinton, "Similarity of neural network representations revisited," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 3519–3529.
- [58] A. X. Chang, T. A. Funkhouser, L. J. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "ShapeNet: An information-rich 3D model repository," 2015, *arXiv:1512.03012*.
- [59] K. Cherenkova, D. Aouada, and G. Gusev, "Pvdeconv: Point-voxel deconvolution for autoencoding CAD construction in 3D," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2020, pp. 2741–2745.
- [60] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3D ShapeNets: A deep representation for volumetric shapes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1912–1920.
- [61] M. A. Uy, Q.-H. Pham, B.-S. Hua, T. Nguyen, and S.-K. Yeung, "Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1588–1597.
- [62] H. Liu, M. Cai, and Y. J. Lee, "Masked discrimination for self-supervised learning on point clouds," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2022.
- [63] R. Zhang, Z. Guo, R. Fang, B. Zhao, D. Wang, Y. Qiao, H. Li, and P. Gao, "Point-M2AE: Multi-scale masked autoencoders for hierarchical point cloud pre-training," 2022, *arXiv:2205.14401*.
- [64] J. Zhou, P. Wang, J. Tang, F. Wang, Q. Liu, H. Li, and R. Jin, "What limits the performance of local self-attention?" *Int. J. Comput. Vis.*, vol. 131, no. 10, pp. 2516–2528, Oct. 2023.
- [65] A. Steiner, A. Kolesnikov, X. Zhai, R. Wightman, J. Uszkoreit, and L. Beyer, "How to train your ViT? Data, augmentation, and regularization in vision transformers," 2021, *arXiv:2106.10270*.
- [66] Q. Zhang, Y. Xu, J. Zhang, and D. Tao, "ViTAEv2: Vision transformer advanced by exploring inductive bias for image recognition and beyond," *Int. J. Comput. Vis.*, vol. 131, no. 5, pp. 1141–1162, May 2023.



IOANNIS ROMANELIS received the Electrical and Computer Engineering Diploma degree from the University of Patras, in 2021, where he is currently pursuing the Ph.D. degree, under the supervision of Prof. Konstantinos Moustakas. He joined the Visualization and Virtual Reality (VVR) Group. His research interests include computer vision, deep learning, point cloud processing, 3D scene understanding, and explainable AI.



VLASSIS FOTIS received the Electrical Engineering and Computer Technology degree from the University of Patras, in 2021, where he is currently pursuing the Ph.D. degree, under the supervision of Prof. Konstantinos Moustakas. He joined the Visualization and Virtual Reality (VVR) Group. His research interests include (but are not limited to) computer vision, theoretical deep learning, 3D scene understanding, and geometry processing.



KONSTANTINOS MOUSTAKAS (Senior Member, IEEE) received the Diploma and Ph.D. degrees in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 2003 and 2007, respectively. From 2007 to 2011, he was a Postdoctoral Research Fellow with the Information Technologies Institute, Centre for Research and Technology Hellas. He is currently a Professor with the Electrical and Computer Engineering Department,

University of Patras; the Head of the Visualization and Virtual Reality Group; and the Director of the Wire Communications and Information Technology Laboratory and the M.Sc. Program on Biomedical Engineering, University of Patras. He serves as an Academic Research Fellow for the ISI/Athena Research Center. During the latest years, he has been the (co)author of more than 300 papers in refereed journals, edited books, and international conferences. He has participated in more than 30 research and development projects funded by the EC and the Greek Secretariat of Research and Technology, while he has served as the Coordinator or the Scientific Coordinator in four of them. His research interests include virtual, augmented and mixed reality, 3D geometry processing, haptics, virtual physiological human modeling, information visualization, physics-based simulations, computational geometry, and computer vision. He has been a member of the organizing committee of several international conferences. He is a Senior Member of the IEEE Computer Society and a member of Eurographics. His research work has received several awards.



ADRIAN MUNTEANU (Member, IEEE) received the M.Sc. degree in electronics and telecommunications from the Politehnica University of Bucharest, Romania, in 1994, the M.Sc. degree in biomedical engineering from the University of Patras, Greece, in 1996, and the Ph.D. degree (maxima cum laude) in applied sciences from Vrije Universiteit Brussel, Belgium, in 2003.

From 2004 to 2010, he was a Postdoctoral Fellow with the Fund for Scientific Research—Flanders (FWO), Belgium. Since 2007, he has been a Professor with Vrije Universiteit Brussel (VUB), Belgium, where he is currently a professor with the Electronics and Informatics (ETRO) Department. He is the author of more than 400 journals and conference publications, book chapters, and contributions to standards. He holds seven patents in image and video coding. His research interests include image, video and 3D graphics compression, 3D video, deep learning, distributed visual processing, error-resilient coding, and multimedia transmission over networks.

Dr. Munteanu was a recipient of the 2004 BARCO-FWO Prize for his Ph.D. work, a (co-)recipient of the Most Cited Paper Award from Elsevier, in 2007, and of ten other scientific prizes and awards. He served as an Associate Editor for IEEE TRANSACTIONS ON MULTIMEDIA. He currently serves as an Associate Editor for IEEE TRANSACTIONS ON IMAGE PROCESSING.

• • •