

M³3D: Learning 3D priors using Multi-Modal Masked Autoencoders for 2D image and video understanding

Muhammad Abdullah Jamal and Omid Mohareri
 Intuitive Surgical Inc., Sunnyvale, CA

Abstract

We present a new pre-training strategy called M³3D (*Multi-Modal Masked 3D*) built based on Multi-modal masked autoencoders that can leverage 3D priors and learned cross-modal representations in RGB-D data. We integrate two major self-supervised learning frameworks; Masked Image Modeling (MIM) and contrastive learning; aiming to effectively embed masked 3D priors and modality complementary features to enhance the correspondence between modalities. In contrast to recent approaches which are either focusing on specific downstream tasks or require multi-view correspondence, we show that our pre-training strategy is ubiquitous, enabling improved representation learning that can transfer into improved performance on various downstream tasks such as video action recognition, video action detection, 2D semantic segmentation and depth estimation. Experiments show that M³3D outperforms the existing state-of-the-art approaches on ScanNet, NYUv2, UCF-101 and OR-AR, particularly with an improvement of +1.3% mIoU against Mask3D on ScanNet semantic segmentation. We further evaluate our method on low-data regime and demonstrate its superior data efficiency compared to current state-of-the-art approaches.

1. Introduction

Pre-training Vision Transformers [24] (ViTs) using Masked Autoencoders (MAE) followed by fine-tuning gives rise to the state-of-the-art results for various computer vision tasks such as image classification [9, 33, 66], video activity recognition [26, 60, 62], semantic segmentation [27, 36, 45] and 3D scene understanding [14, 52, 70]. Inspired by BERT [21], MAE masks high number of patches in an input image or video clip and predicts such missing regions. It usually consists of an asymmetric encoder-decoder architecture, where the encoder objective is to encode the unmasked patches into latent representations while the masked patches are predicted by the decoder using these representations. However, MAEs are mostly focused on learning from single-modality data (images or video clips)

and can't leverage other data modalities that are present in commonly used RGB-D and Time-of-Flight (ToF) cameras. The depth and point cloud data available in such cameras can provide an opportunity to learn geometric priors and avoid view-dependent effects for efficient representation learning. Particularly, various 3D understanding approaches [18, 37, 65, 73] have been leveraging RGB-D datasets for high-level scene understanding and low-level point matching [7, 72] tasks using contrastive learning. However, very little work has been done in exploring 3D priors for 2D image understanding with limited focus on tasks like semantic segmentation and depth estimation.

Our goal is to embed 3D priors into ViTs based backbones for various downstream tasks that include but not limited to video activity recognition, video action detection, semantic segmentation and depth estimation, to effectively learn geometric and structural priors as well as cross-modal features. Recently, Mask3D [36] effectively learns masked 3D priors from a single-view RGB-D frame by only reconstructing the depth through masking out patches in RGB and depth modalities. However, the approach is limited to 2D image understanding tasks and relies on MAE pre-training for any cross-modal representation learning. Pri3D [38] explores to embed the 3D prior by leveraging multi-view RGB-D data to introduce geometric constraints in a contrastive learning scheme for ResNet backbones. However, it relies on camera registration between multiple views for each frame. Instead, we consider to empower ViTs with such 3D priors using MAEs and cross-modal training strategies such as contrastive learning and matching loss that transfer well to various video and image based downstream tasks. Cross-modal learning and Masked Image Modeling are complementary to each other. While the former would leverage the very useful visual-depth pair information to encode modality-invariant features, the latter forces its representation to encode the majority of the input information. This motivates us to integrate both the pre-training strategies to extract different discriminative features.

In this paper, we present M³3D, a multi-modal masked autoencoder based approach which learns to embed 3D priors in a self-supervised manner by pre-training on RGB-D

image pair or video clips and also integrates cross-modal learning. To train M^3D , we randomly mask patches in both RGB and depth modalities and encode the unmasked patches by passing them through either modality-specific transformer encoders or modality-agnostic transformer encoder which is shared between both the modalities. The representations from the encoder are then appended with learned masked tokens to reconstruct masked regions in depth and RGB input. We equipped our approach with contrastive learning and RGB-D matching which are used to enhance cross-modal correspondence during pre-training. The objective of contrastive learning is to align the visual scene and the corresponding depth data by pulling the feature associated with the RGB patch and its corresponding depth while pushing away the other depth patches in the corresponding depth map. Finally, to further enhance the correspondence, the matching loss predicts whether the RGB-depth pair is matched (positive) or not matched (negative) by applying a linear layer followed by softmax on the top of the encoder features to predict a two-class probability. Our experiments demonstrate the effectiveness of M^3D on a variety of datasets for video and image understanding. We pre-train on UCF-101 [57], OR-AR [35] and fine-tune for video activity recognition, and video action detection respectively. We also pre-train the model with ScanNet [20] and fine-tune it for semantic segmentation. We also show the generalizability of the model on NYUv2 [49] data for depth estimation besides semantic segmentation. In summary, the main contributions of our paper are:

- We propose a new self-supervised pre-training approach called M^3D which is equipped with masked autoencoders and cross-modal training for RGB-D representation learning. Our method can be applied to various video and image understanding tasks based on single view RGB-D image pair or video clips.
- Our approach learns to embed 3D priors without any camera pose registration, while enhancing the cross-modal correspondence between RGB and depth modalities using contrastive learning and matching loss.
- We demonstrate the efficacy of M^3D for current ViT backbones on variety of video (UCF-101, OR-AR surgical dataset) and image understanding (ScanNet, NYUv2) datasets.
- We further evaluate our method on low-data regime and demonstrate its superior data efficiency compared to current state-of-the-art approaches. Several ablation studies have been conducted and presented here to prove the effectiveness of various aspects of M^3D .

2. Related Work

Our work proposes a self-supervised learning algorithm for general computer vision tasks. We will review some of the approaches that are closely related to our approach from different aspects - Masked Autoencoder based pre-training for vision transformers, Multi-modal learning and pre-training by learning 3D priors.

Masked Autoencoder based pre-training for transformers. Self-supervised learning paradigm learns visual features based on some pre-text task using large-scale unlabeled dataset, before fine-tuning for multiple downstream tasks. Earlier work on pre-text task includes colorization [71], jigsaw puzzle [50], and predicting image rotation [28]. Instance-discrimination based learning uses two different views or augmentations of an image to learn view-invariant features by pulling them closer in the embedding space while pushing away the negative pair. SimCLR [16], MoCo [69] uses contrastive learning, BYOL [31] uses online and target network while SwAV [11], DeepCluster [10] uses clustering to discriminate between the clusters. Recently, the momentum in self-supervised learning has been shifted towards the masked image modeling which learns useful representations by learning to reconstruct the masked image patches. Inspired by the success in NLP domain such as bidirectional encoder (BERT) [21] and Generative Pre-Training (GPT) [54], several masked image prediction methods for pre-training vision transformers have been proposed with various reconstruction targets such as pixels [4, 15, 24, 25, 33, 66], features [6, 63], discrete tokens via dVAE [9, 74]. They have shown to outperform contrastive learning based approaches on various downstream tasks. One particular approach is MAE [33] which first masks large portion of the input (e.g., 75%) and then passes the unmasked patches to the large encoder followed by a lightweight decoder that reconstruct the masked patches. By masking out large portion of input patches, it accelerates the pre-training of vision transformers. Moreover, these approaches have been extended to the video domain where most of focus is to design a masking strategy that includes random masking [26], tube masking [60] and adaptive masking based on motion [41, 58] or a separate network [8]. On the other-hand, our work proposes multi-modal masked auto-encoder pre-training by learning 3D priors for several downstream tasks such as video action detection, video activity recognition, 2D semantic segmentation and depth estimation.

Multi-Modal Learning. It involves training the models by leveraging data from multiple modalities. The learning may involve training modality-agnostic unified encoder or modality-specific encoders using modalities such as im-

ages and text [2, 13, 19, 39, 42, 59, 68], video and audio [3, 40, 48, 51], video, text and audio [1], and images and depth [29]. One particular approach is MultiMAE [5] which extends MAE to multi-modal multi-task settings and includes depth modality. However, it requires depth during the fine-tuning stage as well as semantic segmentation labels during the pre-training. Our approach doesn't rely on semantic segmentation, it only uses depth during the pre-training and it is not limited to the image domain only. Please refer to experimental section 4 for more results on MultiMAE without semantic segmentation modality.

Pre-training by learning 3D priors. There is a recent surge of learning cross-modal features particularly between languages and images. CLIP [53] learns visual concepts from natural language supervision, showing impressive zero-shot performance for image classification. Pri3D [38] learns view-invariant and geometry-aware representations using multi-view consistency during pre-training for 2D image understanding. It leverages contrastive learning for 2D-3D correspondence and embed 3D priors to pre-train ResNet [34] architectures. However, our approach is more ubiquitous, it is not limited to 2D image understanding tasks and doesn't require camera pose registration across views. Mask3D [36] proposes a masked auto-encoder based self-supervised learning for ViT backbones to learn masked 3D priors for 2D image understanding without the reconstruction of RGB masked patches. In contrast, we formulate a self-supervised pre-training that can operate on both single-view images and videos and leverages masked 3D prior for several downstream tasks. Our approach also enhances the cross-modal learning and 2D-3D correspondence with contrastive learning and matching loss during pre-training which consequently boost the transfer learning performance.

3. Method

We propose M³3D to embed 3D priors, capture modality complementary feature and learn cross-modal representation by self-supervised pre-training using a single-view RGB-D image or a video clip. To learn such priors and representations, we formulate the problem from Masked auto-encoder perspective and cross-modal training strategies.

Given the RGB-D input, we first mask the modalities, and pass them through the modality-specific or modality-agnostic encoder. The latent representations from the encoder(s) are then used to reconstruct the masked RGB and depth input. To learn the cross-modal representations, we use contrastive learning and matching loss which are the two most common cross-modal training strategies. The objective of depth reconstruction task is to embed geometric awareness into the encoder while cross-modal training enhances cross-modal representation learning. After

pre-training, the encoder can be fine-tuned for downstream tasks such video action recognition, 2D semantic segmentation, depth estimation etc. An overview of our approach is shown in Figure 1. Below, we first explain the masking image modeling strategy from multi-modal input perspective. Then, we explain two types of cross-modal training strategies: contrastive, and matching loss. Finally, we describe the overall pre-training pipeline.

3.1. Multi-modal Masked Autoencoder

Our approach is not limited to the images and can be extended to the videos as well. Please see section 4 for performance on the video downstream tasks. For ease of notations, we describe our approach with a single-view RGB-D frame as an input. Given a RGB frame I_R of size $3 \times H \times W$ and its corresponding depth map I_D of size $1 \times H \times W$, we first divide them into $P \times P$ patches from which we randomly keep a percentage of patches by masking out others for both RGB and depth input. $P \times P$ patches are projected to correct dimension using a different linear projection layer for each modality. Please refer to supplementary material on masking strategies. We use Conv2d (Conv3d in case of videos) as a linear projection layer. We then add positional embeddings and pass the tokens correspond to the unmasked patches to the modality-specific encoders or shared encoder between the modalities to get the latent representations. We use ViT-B as an encoder. To reconstruct the masked-out tokens, we use a lightweight shared decoder based on transformers with two prediction head; one for each modality. The input to the decoder is the latent representations of the unmasked patches and the set of mask tokens which are actually placeholders for the decoder to reconstruct the masked patches at the original image resolution. Following [5], we add sine-cosine positional embeddings before passing them to the decoder. From the reconstruction output, we compute the losses only on the masked tokens.

Reconstruction Losses. Our model has two reconstruction loss; one for each modality. Following [33], we first normalize the output patches as well as target patches, and then used compute the MSE loss between the ground truth and the predicted pixels.

$$\mathcal{L}_{rgb} = \frac{1}{\omega} \sum_{p \in \omega} \|I_R(p) - \hat{I}_R(p)\|_2 \quad (1)$$

where p is the token index, ω is the set of masked tokens. \hat{I}_R corresponds to the reconstruction or the prediction of the model.

Similarly, for depth reconstruction task, we follow MultiMAE [5] to use L1 loss for image based pre-training. For

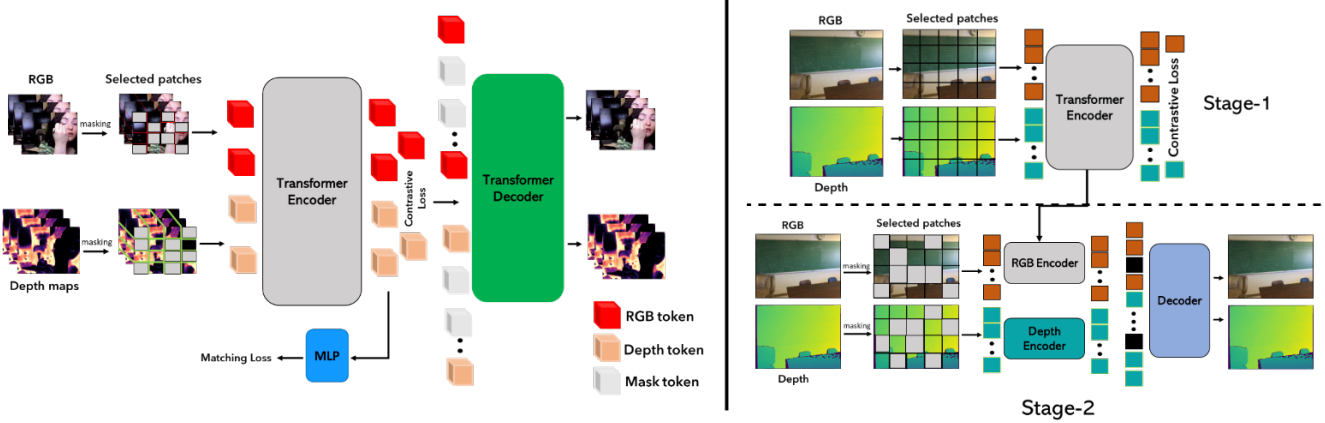


Figure 1. Overview of M³3D pre-training. We introduce two self-supervised learning frameworks; Masked Image Modeling and contrastive learning. For MIM pre-training, we first mask out patches from RGB and depth input and then linearly project the unmasked patches to tokens with a fixed dimension. The tokens are then encoded using a Transformer encoder. The decoder takes the latent representations from the encoder and reconstruct the mask-out patches. Finally, we apply contrastive loss and matching loss on the encoded representations to enable cross-modal correspondence. **(Left)** M³3D pre-training setup for video understanding. **(Right)** M³3D two-stage pre-training setup for image understanding.

video based pre-training, we use MSE Loss.

$$\mathcal{L}_{depth} = \frac{1}{\lambda} \sum_{m \in \lambda} \|I_D(m) - \hat{I}_D(m)\| \quad (2)$$

where m is the token index, λ is the set of masked tokens. \hat{I}_D corresponds to the reconstruction or the prediction of the model.

3.2. Cross-Modal Representation Learning

Besides pre-training our model using masked image modeling, we also propose to use two cross-modal training strategies i.e. RGB-Depth contrastive learning and RGB-Depth matching. The objective of these strategies is to enhance the cross-modal representation learning.

Contrastive Loss. Inspired by the success of contrastive learning [47, 53], we present our cross-modal contrastive pre-training. For each RGB-Depth image pair, we use the latent representations from the transformer encoder. Specifically, we use InfoNCE loss [61] with temperature τ to pre-train the model. The goal of RGB-D contrastive learning is to align the RGB patch and the corresponding depth patch by pulling them closer and repelling unpaired ones apart by minimizing the similarity with the other patches in the corresponding map. Indeed, one can use instance-level contrastive learning, but we hypothesize it will fail to capture high-level information as it is relatively an easy pre-text task.

$$\mathcal{L}_c = -\frac{1}{N} \sum_{i=1}^N \log \left[\frac{\exp(s_{i,i}/\tau)}{\sum_{k \neq i} \exp(s_{i,k}/\tau) + \exp(s_{i,i}/\tau)} \right] \quad (3)$$

where $s_{i,j} = \|X_i^{rgb}\|^T \|X_j^d\|$ and τ is the temperature. $\|X_i^{rgb}\|$ and $\|X_j^d\|$ corresponds to the encoder features for RGB and depth respectively for a patch i .

Matching Loss. Inspired by video-text matching, we also propose to use RGB-Depth matching to further enhance the cross-modal correspondence. It predicts whether the RGB-Depth pair is matched (positive) or not (negative). Specifically, we reuse the features from the encoder, and then passed them to the linear layer followed by a softmax classifier to solve two-class classification problem.

$$\mathcal{L}_m = -\frac{1}{M} \sum_i^M \sum_{t=0}^1 y_{it}^m \log(q_t^m(X_i^{rgb}, X_i^d)) \quad (4)$$

where y_{it}^m is sign function which outputs 1 if value of t is 1 else 0. $q_t^m(X_i^{rgb}, X_i^d)$ denotes the probability score of t from the softmax. M is the total number of RGB-Depth pairs in the batch. We specifically used this loss for video based pre-training.

3.3. Pre-training Setup

The final pre-training loss of M³3D for video based understanding is,

$$\mathcal{L}_{video} = \alpha \mathcal{L}_{rgb} + \beta \mathcal{L}_{depth} + \gamma \mathcal{L}_c + \eta \mathcal{L}_m \quad (5)$$

where $\alpha, \beta, \gamma, \eta$ are the hyper-parameters tuned during pre-training. However, for image based understanding tasks, we pre-train our model in two-stages as shown in Figure 1. In the first stage, we pre-train the encoder using contrastive loss for few epochs and then in the second stage, we initialize the weights of the encoder with stage-1 weights, and then pre-train the encoder and decoder using masked image modeling.

$$\mathcal{L}_{stage1} = \mathcal{L}_c \quad (6)$$

$$\mathcal{L}_{stage2} = \alpha \mathcal{L}_{rgb} + \beta \mathcal{L}_{depth} \quad (7)$$

4. Experiments

M³3D aims to learn masked 3D priors to embed to ViT backbones and enhance the fine-tuning performance with further cross-modal learning using contrastive learning and matching loss. In this section, we first briefly describe the experimental setup for pre-training stage. Then, we provide a transfer study to measure the effectiveness of our approach across various downstream tasks. Moreover, we show the data-efficient nature of our approach by evaluating it under low-data regime setting. Finally, we present some ablation studies to study the effectiveness of each component in our approach.

4.1. Experimental Setup

Pre-train: We use vision-transformer base (ViT-B) as an encoder for all the experiments. For video activity recognition, we pre-train the model on UCF-101 [57] dataset, one of the benchmark datasets for video action recognition. We use monodepth2 [30] to extract the depths from the videos as original dataset contains RGB videos only. Please refer to [30] for more details on the approach. We use a lightweight decoder which consists of 4 blocks following VideoMAE [60]. We initialize the model weights with network weights trained on Kinetics-400 [43]. To maintain the self-supervised paradigm, we initialize the model with weights obtained by self-supervised Kinetics-400 pre-training [60]. We also show the efficacy of our approach by pre-training it from scratch. For surgical video action detection, we pre-train the model on OR-AR [35] dataset which consists of 820 full videos each having 9 temporal workflow phase labels. This dataset is collected using ToF cameras placed in two operating rooms in a single hospital. For 2D image understanding task, we pre-train the model on the ScanNet [20] dataset. ScanNet contains 2.5M RGB-D frames from 1513 video sequences. We regularly sample every 25th frame without any filtering during pre-training. We initialize the encoder weights with the self-supervised ImageNet pre-training [33]. Please refer to the supplementary material for details on datasets, pre-training and hyper-parameters.

4.2. Results

4.2.1 Fine-tuning for video action recognition

We fine-tune M³3D based ViT-B encoder on UCF-101. Table 1 shows the top-1 accuracy comparison with the recent state-of-the-art video based pre-training approaches. It can be seen from the table that our approach achieves 96.3% top-1 accuracy outperforming all the baselines including VideoMAE [60] which shows the effectiveness of leveraging RGB-D data to embed 3D priors for video understanding. Moreover, we also report the results when we pre-train the model from scratch. Although, it achieves the same performance as VideoMAE, M³3D is computationally less expensive than VideoMAE as it requires less epochs and higher masking ratio during pre-training. For fair comparison, we also report the performance of VideoMAE when it is pre-trained for less epochs with higher masking ratio (e.g. 0.90). It is clearly seen from the table that our approach outperforms VideoMAE under the same pre-training epochs and masking ratio. To the best of our knowledge, M³3D provides the first ViT-B backbone that is pre-trained from scratch under self-supervision using RGB-D for video action recognition. It is worth mentioning that the depth maps are extracted from an off-the-shelf model and we hypothesize that the performance will further enhance with an improved depth estimation model or readily available depth data.

4.2.2 Fine-tuning for surgical video action detection

We fine-tune M³3D based ViT-B encoder on OR-AR, a benchmark dataset for surgical video action detection. Table 2 shows the results compared to the recent state-of-the-art video based masked autoencoder pre-training approaches. Following common practice [35, 56], we report the mean average precision (mAP). Our approach consistently outperforms the existing video based MAE models on low-data regime. More specifically, M³3D achieves 80.90 mAP when fine-tuned using 5% labeled data which also shows that it is a more data-efficient learner than VideoMAE. Moreover, we report the results of VideoMAE when it is initialized using Kinetics-400 pre-trained weights. The weights can be found from VideoMAE’s github repository. Similar to video action recognition, it is the first multi-modal based transformer for surgical action detection which shows the effectiveness of the M³3D on datasets in the medical domain. Finally, we report the results of Spatio-temporal MAE [26] with different masking strategies such as random, frame, etc. and they are termed as MAE-Random, MAE-Frame in the table.

Methods	Masking ratio	Backbone	Pre-train Epochs	Dataset	Frames	Top-1
Scratch	-	ViT-B		UCF-101	16	51.4
OPN [44]	-	VGG		UCF-101	N/A	59.6
VCOP [67]	-	R(2+1)D		UCF-101	N/A	72.4
CoCLR [32]	-	S3D-G		UCF-101	32	81.4
Vi ² CLR [22]	-	S3D		UCF-101	32	82.8
CoCLR [32]	-	S3D-G		Kinetics-400	32	87.9
Vi ² CLR [22]	-	S3D		Kinetics-400	32	89.1
MoCov3 [17]	-	ViT-B		UCF-101	16	81.7
VideoMAE [60]		ViT-B	-	Kinetics-400	16	96.0 [†]
M³3D	0.90	ViT-B	100	Kinetics-400 + UCF-101	16	96.3
VideoMAE [60]	0.75	ViT-B	800	UCF-101	16	90.1
VideoMAE [60]	0.90	ViT-B	3200	UCF-101	16	90.7
VideoMAE [60]	0.75	ViT-B	3200	UCF-101	16	91.2
M³3D	0.90	ViT-B	800	UCF-101	16	91.1

Table 1. **Video action recognition accuracy on UCF-101 [57]**. [†] means we re-run the code using the original repository provided by VideoMAE.

Methods	Masking ratio	Backbone	Pre-train	5%	10%	20%	100%
MaskFeat [63]	0.90	MViT-S	Kinetics-400	62.35	78.88	-	-
MAE-Random [26]	0.90	ViT-B	Kinetics-400	64.66	81.48	84.93	94.8
MAE-Random [26]	0.90	ViT-B	OR-AR	66.58	81.87	84.97	96.3
MAE-Frame [26]	0.90	ViT-B	OR-AR	63.44	78.89	81.45	-
VideoMAE [60]	0.90	ViT-B	OR-AR	65.57	81.74	83.89	94.9
VideoMAE [60]	0.90	ViT-B	Kinetics-400 + OR-AR	70.93	83.73	86.33	95.9
M³3D	0.90	ViT-B	Kinetics-400 + OR-AR	80.90	85.31	89.88	96.1

Table 2. **Comparison of M³3D with the other state-of-the-art methods under different data-regime setting on the OR-AR dataset [35]**.

Methods	Reconstruction task	Backbone	Pre-train	Fine-tune Modality	mIoU
Scratch	-	ViT-B	None	RGB	32.6
MultiMAE* [5]	RGB + Depth	ViT-B	ImageNet+ScanNet	RGB	64.0
Pri3D [38]		ViT-B	ImageNet+ScanNet	RGB	59.3
Pri3D [38]	-	ResNet-50	ImageNet+ScanNet	RGB	60.2
DINO [12]	-	ViT-B	ImageNet+ScanNet	RGB	58.1
MAE [33]	RGB	ViT-B	ImageNet	RGB	64.8
MAE [33]	RGB	ViT-B	ImageNet+ScanNet	RGB	64.5
Mask3D [36]	Depth	ViT-B	ImageNet+ScanNet	RGB	66.0
Mask3D [36]	RGB + Depth	ViT-B	ImageNet+ScanNet	RGB	65.5
M³3D	RGB + Depth	ViT-B	ImageNet+ScanNet	RGB	67.3
MultiMAE [5]	RGB + Depth + Segmentation	ViT-B	ImageNet	RGB	66.4

Table 3. **ScanNet Semantic Segmentation**. M³3D outperforms Mask3D and other state-of-the-art approaches that leverage RGB-D data during pre-training.

Methods	Reconstruction task	Backbone	Pre-train	Fine-tune Modality	mIoU
MultiMAE* [5]	RGB + Depth	ViT-B	ImageNet+ScanNet	RGB	49.1
MAE [33]	RGB	ViT-B	ImageNet	RGB	46.9
MAE [33]	RGB	ViT-B	ImageNet+ScanNet	RGB	48.3
Mask3D [36]	Depth	ViT-B	ImageNet+ScanNet	RGB	50.5
M³3D	RGB + Depth	ViT-B	ImageNet+ScanNet	RGB	51.2
MultiMAE [5]	RGB + Depth + Segmentation	ViT-B	ImageNet	RGB	51.5[†]

Table 4. **NYUv2 Semantic Segmentation**. M³3D outperforms state-of-the-art approaches that leverage RGB-D data during pre-training. It demonstrate the effectiveness in transferring to out-domain dataset. [†] means we re-run the code using the original repository provided by MultiMAE.

4.2.3 Fine-tuning for 2D semantic segmentation

In this section, we show that our pre-training approach is capable of transferring the learned representations to the 2D image understanding tasks by performing the 2D semantic segmentation task. Following MultiMAE [5], we use segmentation head based on ConvNext architecture [46] on top of the ViT-B encoder. We use mean Intersection over Union (mIoU) as an evaluation metric. For fine-tuning, we sample every 100th frame, resulting in 20,000 training images and 5000 validation images which is common protocol of the ScanNet benchmark [20]. Table 3 shows the results compared to the recent state-of-the-art pre-training approaches for ScanNet semantic segmentation. It can be observed from the table that M³3D significantly outperforms Mask3D [36] (+1.3 mIoU), a state-of-the-art approach. More notably, our approach improves over MAE [33] which is also pre-trained with ImageNet and ScanNet (+2.8 mIoU). Moreover, compared to Pri3D [38], a 3D-based pre-training method, our approach outperforms by a significant margin of 7.1 mIoU which shows that multi-view based 3D pre-training does not effectively embed 3D priors to ViT backbones, rather degrades the performance compared to ResNet backbone. Furthermore, we also report the results of MultiMAE [5] when it is pre-trained on ScanNet and it is denoted as MultiMAE* in the table. For fair comparison, we only use RGB and depth as the reconstruction task. We observe that M³3D outperforms MultiMAE* by a margin of 3.3 mIoU as well as original MultiMAE. Finally, we demonstrate the generalizability of M³3D across datasets by fine-tuning the pre-trained model on NYUv2 [49] following the same setup. Table 4 shows the results compared to the recent state-of-the-art pre-training approaches for NYUv2 semantic segmentation. We draw the same observation that our pre-trained approach outperforms the competing baselines showing how well it transfers across datasets.

4.2.4 Fine-tuning for depth estimation

In this section, we study how M³3D transfers the representation to the dense regression tasks. We use NYUv2 for the depth estimation and report δ_1 on the NYUv2 test set. δ_1 is the percentage of pixels that have an error ratio ($\max\{\frac{\hat{y}_p}{y_p}, \frac{y_p}{\hat{y}_p}\}$) below 1.25 [23]. Following MultiMAE, we use DPT [55] as dense prediction head on the top of the ViT encoder. Table 5 shows that M³3D outperforms recent state-of-the-art approaches including MultiMAE, Mask3D. Notably, it outperforms CroCo [64] which is a cross view completion pre-training strategy based on MIM specifically designed for 3D vision tasks with 86.7% vs 85.6%. Although our approach is marginally outperforming original MultiMAE, but we want to reiterate that it is pre-trained with three different tasks. Moreover, when multiMAE is




Figure 2. M³3D is a data-efficient learner. We compare to the recent state-of-the-art pre-training approaches on ScanNet 2D semantic segmentation under limited labeled data scenarios. Notably, our approach improves +2.7% mIoU over Mask3D [36] at 20% training data.

pre-trained on ScanNet with RGB and depth only, the performance drops to 85.3%.

4.3. Data-Efficient Learner

We fine-tune M³3D based ViT-B encoder on ScanNet for 2D semantic segmentation mainly under low-data regime setting to study how data-efficient learner our approach is. Figure 2 shows that our approach consistently outperforms the competing approaches by a considerable margin across different percentage of available labeled training data. It is worth mentioning that M³3D recovers more than 80% performance of the full labeled training set performance when fine-tune with only 20% of training data.

4.4. Ablation studies

In this section, we perform in-depth ablation studies on ScanNet 2D semantic segmentation.

Effect of masking ratio during pre-training. We study the influence of masking ratio and report the results in Table 6 on ScanNet semantic segmentation. It is clearly shown from the table that by masking more patches in RGB and depth modality, M³3D achieves the best performance.

Without ImageNet Initialization. We observe a performance drop when the model is not initialized using ImageNet pre-trained weights in the pre-training stage as shown in Table 7. Because ScanNet has a relatively small amount of indoor data, so it’s harder to pre-train the ViT backbones from scratch. Since, ImageNet weights are readily available, so we initialize our models with it following [36, 38].

Methods	Reconstruction task	Backbone	Pre-train	Fine-tune Modality	δ_1
MultiMAE [5]	RGB + Depth	ViT-B	ImageNet+ScanNet	RGB	85.3
MAE [33]	RGB	ViT-B	ImageNet	RGB	85.1
Mask3D [36]	Depth	ViT-B	ImageNet+ScanNet	RGB	85.4
CroCo [64]	RGB + Depth	ViT-B	Habitat	RGB	85.6
MultiMAE* [5]	RGB + Depth + Segmentation	ViT-B	ImageNet	RGB	83.0
M³3D	RGB + Depth	ViT-B	ImageNet+ScanNet	RGB	86.7
MultiMAE [5]	RGB + Depth + Segmentation	ViT-B	ImageNet	RGB	86.4

Table 5. **NYUv2 Depth Estimation.** M³3D outperforms Mask3D and other state-of-the-art approaches that leverage RGB-D data during pre-training. It demonstrate the effectiveness in transferring to dense prediction task for out-domain dataset.* means the number are reported from CroCo [64] paper.

RGB ratio	Depth ratio	mIoU
20.0%	20.0%	65.5
20.0%	50.0%	65.4
20.0%	80.0%	65.4
50.0%	20.0%	66.0
50.0%	50.0%	65.3
80.0%	20.0%	66.1
80.0%	80.0%	67.3

Table 6. We study the effect of different masking ratio for RGB and depth input on ScanNet 2D semantic segmentation where each ratio indicates the percentage of masked patches.

Method	Backbone	mIoU
Train from Scratch	ViT-B	32.6
Mask3D [36]	ViT-B	41.3
M³3D	ViT-B	42.5

Table 7. Results on ScanNet 2D semantic segmentation without ImageNet initialization during pre-training.

Loss	Backbone	mIoU
w/out contrastive loss	ViT-B	65.5
w/out RGB task	ViT-B	66.8
M³3D	ViT-B	67.3

Table 8. Results on on ScanNet 2D semantic segmentation by pre-training without different loss functions.

Effect of each loss function. We also study the effect of loss functions during pre-training and report the performance on ScanNet 2D semantic segmentation. We compare the performance of model without contrastive loss and without RGB reconstruction task. The results are reported in Table 8. From the results, we observe that contrastive loss overall improves the mIoU compared to Mask3D which suggests that cross-modal learning is an important component besides Masked Image Modeling.




Figure 3. **Qualitative Results on ScanNet.** We visualize the predictions of various approaches on 2D Semantic Segmentation task.

5. Conclusion

In this paper, we present M³3D, a new self-supervised pre-training technique with two main functions: (1) learning to embed geometric priors into 2D representations using Masked Image Modeling, (2) learning cross-modal representations in RGB-D data using contrastive learning. M³3D is a general pre-training method applicable to a variety of image/video understanding tasks, doesn't require camera registration between multi-view input as found in recent self-supervised approaches such as Pri3D, and works well in low-data regime. Our extensive experiments on downstream tasks such as video action recognition, video action detection, 2D semantic segmentation and depth estimation show the superiority of M³3D compared to current state-of-the-art 2D representation learning approaches. Future work includes extending the approach to applications with datasets that have more than two modalities.

References

- [1] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text, 2021. [3](#)
- [2] Jean-Baptiste Alayrac, Adrià Recasens, Rosalia Schneider, Relja Arandjelović, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks, 2020. [3](#)
- [3] Relja Arandjelović and Andrew Zisserman. Look, listen and learn, 2017. [3](#)
- [4] Sara Atito, Muhammad Awais, and Josef Kittler. Sit: Self-supervised vision transformer, 2022. [2](#)
- [5] Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. MultiMAE: Multi-modal multi-task masked autoencoders. 2022. [3](#), [6](#), [7](#), [8](#)
- [6] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. data2vec: A general framework for self-supervised learning in speech, vision and language, 2022. [2](#)
- [7] Mohamed El Banani and Justin Johnson. Bootstrap your own correspondences, 2021. [1](#)
- [8] Wele Gedara Chaminda Bandara, Naman Patel, Ali Gholami, Mehdi Nikkhah, Motilal Agrawal, and Vishal M. Patel. Adamae: Adaptive masking for efficient spatiotemporal learning with masked autoencoders, 2022. [2](#)
- [9] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEit: BERT pre-training of image transformers. In *International Conference on Learning Representations*, 2022. [1](#), [2](#)
- [10] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features, 2019. [2](#)
- [11] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems*, volume 33, pages 9912–9924. Curran Associates, Inc., 2020. [2](#)
- [12] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers, 2021. [6](#)
- [13] Lluís Castrejon, Yusuf Aytar, Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Learning aligned cross-modal representations from weakly aligned data, 2016. [3](#)
- [14] Anthony Chen, Kevin Zhang, Renrui Zhang, Zihan Wang, Yuheng Lu, Yandong Guo, and Shanghang Zhang. Pimae: Point cloud and image interactive masked autoencoders for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. [1](#)
- [15] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research. PMLR, 2020. [2](#)
- [16] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. [2](#)
- [17] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9640–9649, October 2021. [6](#)
- [18] Yujin Chen, Matthias Nießner, and Angela Dai. 4dcontrast: Contrastive learning with dynamic correspondences for 3d scene understanding, 2022. [1](#)
- [19] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning, 2020. [3](#)
- [20] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2017. [2](#), [5](#), [7](#)
- [21] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, 2019. [1](#), [2](#)
- [22] Ali Diba, Vivek Sharma, Reza Safdari, Dariush Lotfi, Saquib Sarfraz, Rainer Stiefelhagen, and Luc Van Gool. Vi2clr: Video and image for visual contrastive learning of representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1502–1512, October 2021. [6](#)
- [23] Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning, 2017. [7](#)
- [24] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [1](#), [2](#)
- [25] Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jégou, and Edouard Grave. Are large-scale datasets necessary for self-supervised pre-training?, 2021. [2](#)
- [26] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *arXiv:2205.09113*, 2022. [1](#), [2](#), [5](#), [6](#)
- [27] Peng Gao, Teli Ma, Hongsheng Li, Jifeng Dai, and Yu Qiao. Convmae: Masked convolution meets masked autoencoders. *arXiv preprint arXiv:2205.03892*, 2022. [1](#)
- [28] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Un-supervised representation learning by predicting image rotations, 2018. [2](#)
- [29] Rohit Girdhar, Mannat Singh, Nikhila Ravi, Laurens van der Maaten, Armand Joulin, and Ishan Misra. Omnivore: A single model for many visual modalities, 2022. [3](#)

- [30] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth prediction. October 2019. [5](#)
- [31] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, volume 33, 2020. [2](#)
- [32] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. In *Neurips*, 2020. [6](#)
- [33] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [34] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [3](#)
- [35] Zhuohong He, Ali Mottaghi, Aidean Sharghi, Muhammad Abdullah Jamal, and Omid Mohareri. An empirical study on activity recognition in long surgical videos. In *Proceedings of the 2nd Machine Learning for Health symposium*, Proceedings of Machine Learning Research. PMLR, 2022. [2](#), [5](#), [6](#)
- [36] Ji Hou, Xiaoliang Dai, Zijian He, Angela Dai, and Matthias Nießner. Mask3d: Pre-training 2d vision transformers by learning masked 3d priors, 2023. [1](#), [3](#), [6](#), [7](#), [8](#)
- [37] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts, 2021. [1](#)
- [38] Ji Hou, Saining Xie, Benjamin Graham, Angela Dai, and Matthias Nießner. Pri3d: Can 3d priors help 2d representation learning?, 2021. [1](#), [3](#), [6](#), [7](#)
- [39] Ronghang Hu and Amanpreet Singh. Unit: Multimodal multitask learning with a unified transformer, 2021. [3](#)
- [40] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppala, Daniel Zoran, Andrew Brock, Evan Shelhamer, Olivier J Henaff, Matthew Botvinick, Andrew Zisserman, Oriol Vinyals, and Joao Carreira. Perceiver IO: A general architecture for structured inputs & outputs. In *International Conference on Learning Representations*, 2022. [3](#)
- [41] Muhammad Abdullah Jamal and Omid Mohareri. Surgmae: Masked autoencoders for long surgical video analysis, 2023. [2](#)
- [42] Lukasz Kaiser, Aidan N. Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. One model to learn them all, 2017. [3](#)
- [43] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *CoRR*, abs/1705.06950, 2017. [5](#)
- [44] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Kumar Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequence. In *IEEE International Conference on Computer Vision*, 2017. [6](#)
- [45] Gang Li, Heliang Zheng, Daqing Liu, Chaoyue Wang, Bing Su, and Changwen Zheng. Semmae: Semantic-guided masking for learning masked autoencoders, 2022. [1](#)
- [46] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s, 2022. [7](#)
- [47] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos, 2020. [4](#)
- [48] Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. Attention bottlenecks for multimodal fusion, 2022. [3](#)
- [49] Pushmeet Kohli, Nathan Silberman, Derek Hoiem, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. [2](#), [7](#)
- [50] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles, 2017. [2](#)
- [51] Andrew Owens and Alexei A. Efros. Audio-visual scene analysis with self-supervised multisensory features, 2018. [3](#)
- [52] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 604–621. Springer, 2022. [1](#)
- [53] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. [3](#), [4](#)
- [54] Alec Radford and Karthik Narasimhan. Improving language understanding by generative pre-training. 2018. [2](#)
- [55] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *ArXiv preprint*, 2021. [7](#)
- [56] Aidean Sharghi, Helene Haugerud, Daniel Oh, and Omid Mohareri. Automatic operating room surgical activity recognition for robot-assisted surgery. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, 2020. [5](#)
- [57] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [2](#), [5](#), [6](#)
- [58] Xinyu Sun, Peihao Chen, Liangwei Chen, Changhao Li, Thomas H Li, Mingkui Tan, and Chuang Gan. Masked motion encoding for self-supervised video representation learning. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#)
- [59] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers, 2019. [3](#)
- [60] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *Advances in Neural*

Information Processing Systems, volume 35, pages 10078–10093. Curran Associates, Inc., 2022. [1](#), [2](#), [5](#), [6](#)

- [61] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2019. [4](#)
- [62] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. Videomae v2: Scaling video masked autoencoders with dual masking, 2023. [1](#)
- [63] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *CVPR*, 2022. [2](#), [6](#)
- [64] Weinzaepfel, Philippe and Leroy, Vincent and Lucas, Thomas and Brégier, Romain and Cabon, Yohann and Arora, Vaibhav and Antsfeld, Leonid and Chidlovskii, Boris and Csurka, Gabriela and Revaud Jérôme. CroCo: Self-Supervised Pre-training for 3D Vision Tasks by Cross-View Completion. In *NeurIPS*, 2022. [7](#), [8](#)
- [65] Saining Xie, Jiatao Gu, Demi Guo, Charles R. Qi, Leonidas J. Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding, 2020. [1](#)
- [66] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [2](#)
- [67] Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yueting Zhuang. Self-supervised spatiotemporal learning via video clip order prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [6](#)
- [68] Haiyang Xu, Ming Yan, Chenliang Li, Bin Bi, Songfang Huang, Wenming Xiao, and Fei Huang. E2e-vlp: End-to-end vision-language pre-training enhanced by visual learning, 2021. [3](#)
- [69] Chaoning Zhang, Kang Zhang, Trung X. Pham, Axi Niu, Zhinan Qiao, Chang D. Yoo, and In So Kweon. Dual temperature helps contrastive learning without many negative samples: Towards understanding and simplifying moco, 2022. [2](#)
- [70] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Pointm2ae: Multi-scale masked autoencoders for hierarchical point cloud pre-training. *arXiv preprint arXiv:2205.14401*, 2022. [1](#)
- [71] Richard Zhang, Phillip Isola, and Alexei A. Efros. Colorful image colorization, 2016. [2](#)
- [72] Yu Zhang, Junle Yu, Xiaolin Huang, Wenhui Zhou, and Ji Hou. Pcr-cg: Point cloud registration via deep color and geometry, 2023. [1](#)
- [73] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-supervised pretraining of 3d features on any point-cloud, 2021. [1](#)
- [74] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer, 2022. [2](#)