

# MILAN: MASKED Image PRETRAINING ON LANGUAGE ASSISTED REPRESENTATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Self-attention based transformer models have been dominating many computer vision tasks in the past few years. Their superb model qualities heavily depend on the excessively large labeled image datasets. In order to reduce the reliance on large labeled datasets, reconstruction based masked autoencoders are gaining popularity, which learn high quality transferable representations from unlabeled images. For the same purpose, recent weakly supervised image pretraining methods explore language supervision from text captions accompanying the images. In this work, we propose masked image pretraining on language assisted representation, dubbed as MILAN. Instead of predicting raw pixels or low level features, our pretraining objective is to reconstruct the image features with substantial semantic signals that are obtained using caption supervision. Moreover, to accommodate our reconstruction target, we propose a more efficient prompting decoder architecture and a semantic aware mask sampling mechanism, which further advance the transfer performance of the pretrained model. Experimental results demonstrate that MILAN delivers higher accuracy than the previous works. When the masked autoencoder is pretrained and finetuned on ImageNet-1K dataset with an input resolution of  $224 \times 224$ , MILAN achieves a top-1 accuracy of 85.4% on ViT-Base, surpassing previous state-of-the-arts by 1%. In the downstream semantic segmentation task, MILAN achieves 52.7 mIoU using ViT-Base backbone on ADE20K dataset, outperforming previous masked pretraining results by 4 points.

## 1 INTRODUCTION

In recent years, we have seen a wide adoption of applying natural language processing (NLP) techniques in computer vision (CV) tasks. The vision transformer (ViT) model (Dosovitskiy et al., 2020) applies the self-attention based transformer architecture to vision tasks and have achieved remarkable performance. However, training ViT models requires much larger labeled datasets to avoid overfitting, such as ImageNet-22K (Deng et al., 2009) and JFT-300M (Sun et al., 2017). Explicitly labeling large image datasets is hardly affordable.

Reconstruction based self-supervised pretraining can extract semantic information from unlabeled data, and has become a popular method to reduce the reliance on very large labeled datasets in both NLP and CV. It is first exemplified by BERT (Devlin et al., 2018) in NLP. Acting like a masked autoencoder (Vincent et al., 2010), BERT randomly masks some percentage of the input word tokens and learns to reconstruct the vocabularies of those masked tokens. Works in (Bao et al., 2021; He et al., 2021; Xie et al., 2021b; Wei et al., 2021) adopt similar techniques in CV to address the data-hungry issue of ViT models. A large percentage of the input image patches are randomly masked with the goal of reconstructing them.

The mask autoencoders can be extended in several directions. Unlike the masked word tokens in NLP, which contain rich semantic information, the masked image patches only contain low-level pixel data. Several works (Bao et al., 2021; Dong et al., 2021; Chen et al., 2022) explore more abstract **reconstruction targets**, aiming to learn higher level visual concepts. However, those methods still only retrieve semantic signals from raw image pixels, which by itself is a difficult task. In addition, the selection of the reconstruction targets heavily influences the **decoder design** in autoencoders, as the decoder serves to reconstruct the masked features with the guidance from the encoder’s output representations. Full fledged transformer blocks are used in the decoder of MAE (He et al., 2021) to

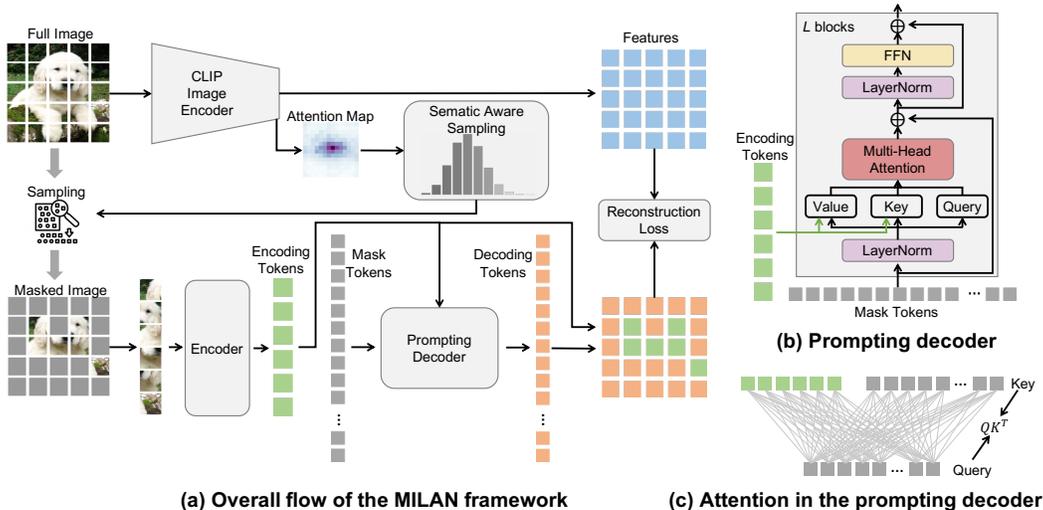


Figure 1: (a) The overall flow of MILAN. The masked autoencoder uses the outputs of the CLIP model as the reconstruction target. An efficient prompting decoder freezes the features of the encoding tokens and only updates the mask tokens. A semantic aware sampling is used to guide the selection of the unmasked image patches. The reconstruction loss is computed on the representation features of both masked and unmasked patches. (b) A detailed diagram of the prompting decoder. (c) The attention computation in the prompting decoder.

reconstruct masked input patches pixel by pixel, whereas lightweight linear layer is adopted in the decoder of MaskFeat (Wei et al., 2021) to reconstruct local features of the image. Thus, if we were using more semantic preserving reconstruction targets, a task tailored decoder architecture would be required. Furthermore, different **sampling strategies** (e.g., grid, block, random) of the input image patches affect the final performance of masked image pretraining (He et al., 2021; Xie et al., 2021b). Majority of prior arts (Bao et al., 2021; Dong et al., 2021; He et al., 2021; Xie et al., 2021b; Wei et al., 2021; Baeviski et al., 2022; Chen et al., 2022) sample the masked patches uniformly at random since it is unbiased and can guarantee coverage. However, it is indifferent to more discriminative image patches and unimportant ones, thus may suffer from slow training convergence (Kakogeorgiou et al., 2022).

In this work, we analyze three highly correlated aspects in masked autoencoders: the reconstruction target, the decoder design, and the mask sampling strategy. We propose a new approach called MILAN, which performs masked image pretraining on language assisted representations. In specific:

(1) We recognize the limitation of extracting semantic signals from raw image pixels alone. But such signals are readily available in the captions accompanying the images. Recent works such as CLIP (Radford et al., 2021), SLIP (Mu et al., 2021), and COCA (Yu et al., 2022) explore the use of caption supervision to learn image representations on abundant image-text pairs obtained from the Internet. The output image features from those models implicitly contain semantic information that facilitates the interpretation of the image contents. In this work, we take the image features coming out of the CLIP image encoder (Radford et al., 2021) as the reconstruction targets for the masked image pretraining, which benefits from natural language supervision and encourages the model to learn high level visual concepts. More interestingly, we will show that the quality of the representation improves on the targets after masked image pretraining.

(2) We realize the tight coupling between the decoder architecture and the reconstruction targets. We design an efficient prompting decoder suitable for reconstruction targets that are latent representations containing affluent semantic signals. It freezes the encoder’s output representations of the unmasked patches and uses them as “fixed prompts” to reconstruct the features of the masked patches. Prompting decoder achieves higher accuracy and reduces the decoding computational cost simultaneously.

(3) Different image patch sampling strategies impact pretraining efficiency. Since our reconstruction targets provide global structure information of the images, we propose a semantic aware mask

sampling mechanism to discriminate semantically important image patches from the insignificant background patches, which improves representation quality and pretraining efficiency.

(4) Combining the three aspects leads to our MILAN framework (Figure 1). Experimentally, our ViT-Base and ViT-Large models pretrained and finetuned on ImageNet-1K dataset achieve 86.4% and 88.3% top-1 accuracy, respectively. Moreover, MILAN significantly boosts the linear probing accuracy compared to reconstruction based and language-image based pretraining methods, and achieves state-of-the-art performance on the downstream object detection, instance segmentation, and semantic segmentation tasks.

## 2 METHODOLOGY

### 2.1 OVERVIEW

The overall flow of MILAN is illustrated in Figure 1(a). We use a masked autoencoder architecture similar to MAE (He et al., 2021). The encoder transforms the unmasked patches into latent representations. The decoder reconstructs the representations of the masked patches assisted by the features of the unmasked patches. We use the latent features that the CLIP image encoder produces from the full image as the reconstruction targets, which are contextualized representations based on the global structure of the image and caption information. The attention map is extracted from the last self-attention layer of the CLIP image encoder and is used to construct a semantic aware sampling distribution to sample unmasked patches. The sampled patches are sent into the encoder and mapped to the latent feature space. We design a prompting decoder that freezes the encoder’s output when hallucinating the features of the masked patches from mask tokens. As shown in Figure 1(b), the query of the attention block in the prompting decoder only contains the features of mask tokens. The key and value matrices comprise both the encoder’s outputs and the features of mask tokens. The masked patches’ features from the prompting decoder and the unmasked patches’ features from the encoder are combined at the end. We re-order the combined full set of features to align with the targets, and compute the reconstruction loss.

MILAN differs from its closely related MAE in: 1) the targets we predict are latent representations obtained with language guidance, whereas MAE reconstructs raw pixels; 2) mask sampling in MILAN is more adapted to patches’ discriminativeness in contrast to MAE’s uniform sampling; 3) our prompting decoder does not update the encoder’s output and thus is more efficient.

### 2.2 RECONSTRUCTION TARGET: LANGUAGE ASSISTED REPRESENTATION

The reconstruction target is a crucial component in masked image pretraining. It influences the semantics of the learned latent representations. Language naturally contains rich semantics, while such information is more difficult to extract from image pixels directly. Thus, an image-text pair provides more meaningful learning signals than an image alone. In practice, texts accompanying images can be easily obtained at scale in the form of image captions. Such large unlabeled image-text datasets foster a series of weakly-supervised image pretraining methods (Radford et al., 2021; Mu et al., 2021; Li et al., 2021a; Singh et al., 2022) with caption supervision. It is expected that the visual representations learned with language guidance can provide affluent semantic information. Therefore, we take the language assisted representations as the reconstruction target in our masked image pretraining framework.

In this work, we primarily use the pretrained CLIP (Radford et al., 2021) model to generate the reconstruction targets. CLIP is trained on a dataset containing images and free-form text captions with InfoNCE loss (Van den Oord et al., 2018). The model has an image encoder and a text encoder, both using the transformer architecture. The encoded image and text features are projected to the same dimension and normalized. Given a batch of image-text pairs, CLIP trains the image encoder and the text encoder by maximizing the feature cosine similarity for the matching image-text pairs in the batch while minimizing the cosine similarity for all other non-matching pairs. Though without labeling data, the image features are trained to be close to the paired text features and distill the rich semantics embedded in the text features. This is illustrated in Figure 2, where the last layer features produced by MAE’s pretrained model, CLIP’s image encoder, and our pretrained model are visualized by t-SNE plots. As shown, the representations from CLIP’s image encoder for each category tend to be grouped together while the pretrained model from MAE cannot distinguish the visual concepts in

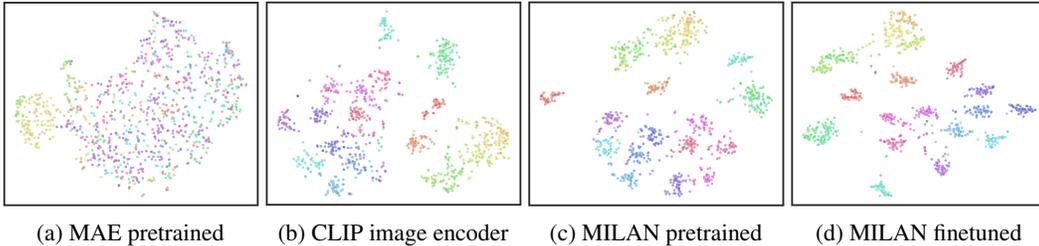


Figure 2: t-SNE visualization of the learned features from ViT-Base obtained by different pretraining methods. We plot the features before the final linear head. We use images of randomly sampled 20 classes in ImageNet-1K validation split.

different categories. MILAN adopts the image features from CLIP as the target and trains the model with a more challenging masked prediction objective. The learned representations are better clustered for different categories.

The pretraining objective of MILAN is formally described as follows. Let  $f_\theta$  denote the CLIP image encoder, whose weights  $\theta$  are frozen. The masked autoencoder under training comprises encoder  $g_\xi$  with weights  $\xi$  and decoder  $h_\nu$  with weights  $\nu$ . From a given full image  $\mathbf{x}$ , the CLIP image encoder outputs the target features  $\{\mathbf{t}_j\}_{j=1}^N = f_\theta(\mathbf{x})$  where  $N$  is the number of image patches. We mask a high portion of patches in  $\mathbf{x}$ , and obtain a masked image  $\tilde{\mathbf{x}}$ . The masked autoencoder outputs the reconstructed representations  $\{\mathbf{p}_j\}_{j=1}^N = (h_\nu \circ g_\xi)(\tilde{\mathbf{x}})$ . We apply  $\ell_2$ -normalization to both the targets and reconstructions:  $\bar{\mathbf{t}}_j = \mathbf{t}_j / \|\mathbf{t}_j\|_2$ ,  $\bar{\mathbf{p}}_j = \mathbf{p}_j / \|\mathbf{p}_j\|_2$  for  $j \in [N]$ . Finally, we define the following mean squared error between the normalized target features and reconstructed representations:

$$\mathcal{L}_{\xi, \nu} = (1/N) \cdot \sum_{j=1}^N \|\bar{\mathbf{p}}_j - \bar{\mathbf{t}}_j\|_2^2. \quad (1)$$

MILAN learns the weights  $\xi, \nu$  of the masked autoencoder by minimizing objective 1. Note that the reconstruction loss is computed on the features of both masked and unmasked patches.

### 2.3 DECODER DESIGN: PROMPTING DECODER

Since the decoder is discarded after pretraining, the encoder needs to learn rich semantic information in the latent representations of the unmasked patches from the reconstruction target. To achieve so, the functional roles of the encoder and decoder need to be clearly segregated. All representation learning for the unmasked patches is completed in the encoder, while the decoder is only for predicting the target features of the masked patches. However, in some previous works (Xie et al., 2021b; Bao et al., 2021), the decoder is as simple as one linear layer, which may be insufficient to reconstruct the masked representations, and portions of the encoder may serve as the decoder. MAE (He et al., 2021) uses a deep decoder that not only updates the mask tokens but also enhances the features of the unmasked patches. Because our method reconstructs language guided latent representations instead of raw pixels, the encoder’s outputs should only provide clues to the decoder to complete the missing patches’ features without being updated in the decoder. Otherwise, the representation quality from the encoder becomes sub-optimal.

In MILAN, we propose a prompting decoder shown in Figure 1(b), where the representations of the unmasked patches from the encoder are frozen, serving as “fixed prompts”. They are appended to the keys and values in each attention module of the prompting decoder’s transformer block while the queries only contain the features of mask tokens. In specific, the multi-head attention (MHA) module in Figure 1(b) performs the following operations:

$$\begin{aligned} \text{MHA}(X, Z) &= \text{Attn}(XW_q, \text{concat}(Z, XW_k), \text{concat}(Z, XW_v)), \\ &= \sum_{h=1}^H \text{softmax}\left(\frac{XW_q^h \text{concat}(Z, XW_k^h)^T}{\sqrt{d^h}}\right) \text{concat}(Z, XW_v^h) W_o^h. \end{aligned} \quad (2)$$

In equation 2,  $X$  and  $Z$  are the features of mask tokens and the encoder’s outputs, respectively. “Attn” is short for the softmax attention operation. “concat” means concatenating along the sequence dimension.  $H$  is the number of heads.  $W_q^h, W_k^h, W_v^h, W_o^h$  represent the query, key, value and output projection weight matrices in each head.  $d^h$  is the embedding dimension in each head. As shown in Figure 1(c), our prompting decoder only computes the self-attention among the features of

mask tokens and the cross-attention between the encoder’s output and the features of mask tokens. Moreover, the FFN modules in the prompting decoder only compute on the features of mask tokens.

Using the default 75% masking ratio, our prompting decoder reduces the decoding computation cost by 20% compared to MAE (He et al., 2021). More importantly, prompting decoder improves the finetuning accuracy significantly, as will be analyzed in Section 3.1.

## 2.4 MASKING STRATEGY: SEMANTIC AWARE SAMPLING

To make the masked image pretraining a meaningful pretext task, previous works mask a very high portion of input image patches uniformly at random. With this aggressive masking strategy, it is possible that the remaining few visible patches only contain background information, which may not provide the important clues needed to reconstruct the foreground objects buried in the piles of masked patches, obstructing the model to learn transferable representations. This becomes a more severe problem in our framework, because the latent representations from the encoder are frozen in the decoding process. To ensure the representation quality of the pretrained model, previous methods (He et al., 2021; Baevski et al., 2022; Wei et al., 2021) usually require very long pretraining epochs.

To improve the pretraining efficiency, we propose a semantic aware mask sampling strategy that can make more rational decisions on which patches to mask when using a very high masking ratio. The idea is that the few visible patches fed into the encoder cover important image regions with high probabilities, so that the latent representations from the encoder provide sufficient clues to the decoder to predict the representations of the masked patches.

To discriminate the semantically important patches from the unimportant ones, we use the attention weights from the last self-attention layer in the CLIP image encoder, which takes the patches of the entire image and an extra class token as input. Denote the input features to the last self-attention layer of the CLIP image encoder by  $[\mathbf{z}_{\text{class}}; \mathbf{z}_1; \dots; \mathbf{z}_N] \in \mathbb{R}^{(N+1) \times d}$ , where  $N$  is the sequence length and  $d$  is the embedding dimension. The interaction between the class token and other features is given by the following attention mechanism:

$$\mathbf{s}_{\text{class}} = \text{softmax}(\mathbf{q}_{\text{class}} \mathbf{K}^T / \sqrt{d}), \quad (3)$$

where  $\mathbf{s}_{\text{class}} \in \mathbb{R}^{1 \times (1+N)}$  is the attention vector of the class token.  $\mathbf{q}_{\text{class}} = \mathbf{z}_{\text{class}} W_q$  is the query vector of the class token, and  $\mathbf{K} = [\mathbf{z}_{\text{class}}; \mathbf{z}_1; \dots; \mathbf{z}_N] W_k$  is the key matrix, where the query and key projection matrices have dimensions  $W_q, W_k \in \mathbb{R}^{d \times d}$ . For simplicity, we show a single-head attention in equation 3. When multiple attention heads are present,  $\mathbf{s}_{\text{class}}$  is obtained by averaging over all the heads. Because the class token from the last layer of the CLIP image encoder is used to align with the text embedding from the text encoder,  $\mathbf{s}_{\text{class}}$  reflects how much information one image patch contributes to the output features of the CLIP image encoder. The magnitude of the  $i$ -th element in  $\mathbf{s}_{\text{class}}$ , denoted by  $\mathbf{s}_{\text{class}}(i)$ , indicates whether the  $i$ -th patch is semantically important or not. The attention vector  $\mathbf{s}_{\text{class}}$  provides us the premise to design a non-uniform sampling distribution. Due to the softmax operation, we can regard  $\mathbf{s}_{\text{class}}(i)$  as the probability of leaving the  $i$ -th patch unmasked in the input image. Let  $r$  represent the masking ratio. The indices of the unmasked patches are obtained by sampling a Multinomial distribution according to the probabilities  $\{\mathbf{s}_{\text{class}}(0), \dots, \mathbf{s}_{\text{class}}(N)\}$  for  $\lceil (1-r)N \rceil$  trials without replacement.

## 3 EXPERIMENTS

We pretrain the ViT-Base and ViT-Large models using MILAN method on ImageNet-1K dataset for 400 epochs using PyTorch framework on A100 machines. The detailed training setup and hyperparameters can be found in the appendix. We use the CLIP ViT-Base and the CLIP ViT-Large image encoders obtained from OpenAI’s paper (Radford et al., 2021) to produce the reconstruction targets when pretraining our ViT-Base and ViT-Large models, respectively.

### 3.1 CLASSIFICATION ON IMAGENET-1K

**Finetuning results.** Table 1 compares the finetuning accuracy on ImageNet-1K dataset using MILAN and previous works on the ViT model architecture. We pretrain and finetune the ViT models using ImageNet-1K dataset only. Since the CLIP model we use is pretrained on OpenAI’s in-house 400M data, we also list it in the training data for MILAN. However, we only use its image encoder’s output features as the reconstruction target for our masked autoencoder. Even though the supervised ViT models are pretrained on the large JFT300M dataset with explicit human labels, MILAN still outperforms them by a clear margin, *e.g.*, improving ViT-Base by +2.2%.

Method	Training data	Res.	ViT-Base		ViT-Large	
			Epochs	Top-1 (%)	Epochs	Top-1 (%)
Supervised (Touvron et al., 2022)	IN1K	224	-	83.8 (+1.6)	-	84.9 (+2.9)
<i>contrastive or clustering based</i>						
MoCov3 (Chen et al., 2021)	IN1K	224	300	83.2 (+2.2)	300	84.1 (+3.7)
DINO (Caron et al., 2021)	IN1K	224	400	82.8 (+2.6)	-	-
iBOT (Zhou et al., 2021)	IN22K&1K	224	320	84.4 (+1.0)	200	86.3 (+1.5)
<i>reconstruction based</i>						
BEiT (Bao et al., 2021)	D250M+IN22K&1K	224	150	83.7 (+1.7)	150	86.0 (+1.8)
mc-BEiT (Li et al., 2022)	OI9M+IN1K	224	800	84.1 (+1.3)	800	85.6 (+2.2)
PeCo (Dong et al., 2021)	IN1K	224	800	84.5 (+0.9)	800	86.5 (+1.3)
SimMIM (Xie et al., 2021b)	IN1K	224	800	83.8 (+1.6)	-	-
MaskFeat (Wei et al., 2021)	IN1K	224	1600	84.0 (+1.4)	1600	85.7 (+2.1)
data2vec (Baevski et al., 2022)	IN1K	224	800	84.2 (+1.2)	1600	86.6 (+1.2)
CAE (Chen et al., 2022)	IN1K	224	800	83.6 (+1.8)	-	-
MAE (He et al., 2021)	IN1K	224	1600	83.6 (+1.8)	1600	85.9 (+1.9)
<i>language-image pretraining based</i>						
CLIP (Radford et al., 2021)	OpenAI400M+IN1K	224	-	82.1 (+3.3)	-	85.3 (+2.5)
MVP (Wei et al., 2022)	OpenAI400M+IN1K	224	300	84.4 (+1.0)	300	86.3 (+1.5)
<b>MILAN</b>	OpenAI400M+IN1K	224	400	<b>85.4</b>	400	<b>87.8</b>
Supervised (Dosovitskiy et al., 2020)	JFT300M+IN1K	384	90	84.2 (+2.2)	90	87.1 (+1.2)
BEiT (Bao et al., 2021)	D250M+IN1K	384	800	84.6 (+1.8)	800	86.3 (+2.0)
SWAG (Singh et al., 2022)	IG3.6B+IN1K	384	2	85.3 (+1.1)	-	-
<b>MILAN</b>	OpenAI400M+IN1K	384	400	<b>86.4</b>	400	<b>88.3</b>

Table 1: Comparison of the **finetuning** top-1 accuracy on ImageNet-1K dataset. All models are pretrained with  $224 \times 224$  input resolution. We compare finetuning with both  $224 \times 224$  and  $384 \times 384$  resolutions. “Epochs” refer to the pretraining epochs. “-”: not reported by the original paper.

The self-supervised pretraining methods are divided by using contrastive or reconstruction based objectives. We also compare with large-scale weakly-supervised pretraining using hashtag supervision (Singh et al., 2022) from an external dataset of 3.6 billion training samples. MILAN produces higher accuracy than all listed prior arts. Compared with MAE, MILAN improves the accuracy by +1.8% for ViT-Base and +1.9% for ViT-Large.

The CLIP model learns visual representations with language supervision on a large image-text dataset. Finetuning the CLIP image encoder does not lead to competitive accuracy. However, when using the image features from the pretrained CLIP model as the reconstruction target to train a mask autoencoder, MILAN improves the accuracy by 3.3% on ViT-Base.

**Linear probing results.** Instead of finetuning the entire model, we also perform linear probing by appending a linear classifier after the final layer of the pretrained model, and only finetune the linear classifier. Table 2 compares the top-1 accuracy on ImageNet-1K dataset using various pretraining methods. As the results show, MILAN beats the accuracy of reconstruction based and language-image pretraining based approaches by a large margin. It matches (on ViT-Base) and outperforms (on ViT-Large) previous best contrastive based methods, which learn more linearly separable representations by instance discrimination (Chen et al., 2021) or clustering (Caron et al., 2021), and are known to be more effective in linear probing (Chen et al., 2020b; He et al., 2020; Chen & He, 2021; Grill et al., 2020). More linear probing results by other variants of MILAN can be found in the appendix.

### 3.2 DOWNSTREAM TASKS

**Object detection and instance segmentation on COCO.** To verify the transferability of MILAN, we evaluate it on COCO dataset (Lin et al., 2014) for object detection and instance segmentation. Following MAE (He et al., 2021), the pretrained ViT backbones are adapted to FPN (Lin et al., 2017) in the Mask R-CNN framework (He et al., 2017), which is finetuned end-to-end on COCO training set to produce the bounding boxes (evaluated by box AP) and the instance masks (evaluated by mask AP) simultaneously. The results are shown in Table 3. Compared to supervised pretraining, MILAN performs better in both tasks, achieving 4.7 and 2.6 points improvements by  $AP_{\text{box}}$  and  $AP_{\text{mask}}$  on ViT-Base, respectively. Compared with the previous best result from MAE, which is obtained by

Method	ViT-Base		ViT-Large	
	Epochs	Top-1 (%)	Epochs	Top-1 (%)
<i>contrastive or clustering based</i>				
MoCov3 (Chen et al., 2021)	300	76.7 (+3.2)	300	77.6 (+6.7)
DINO (Caron et al., 2021)	400	78.2 (+1.7)	-	-
iBoT (Zhou et al., 2021)	1600	79.5 (+0.4)	1000	81.0 (+3.3)
<i>reconstruction based</i>				
BEiT (Bao et al., 2021)	800	56.7 (+23.2)	800	73.5 (+10.8)
SimMIM (Xie et al., 2021b)	800	56.7 (+23.2)	-	-
MaskFeat (Wei et al., 2021)	-	-	1600	67.7 (+16.6)
CAE (Chen et al., 2022)	800	68.3 (+11.6)	-	-
MAE (He et al., 2021)	1600	68.0 (+11.9)	1600	75.8 (+8.5)
<i>language-image pretraining based</i>				
CLIP (Radford et al., 2021)	-	66.5 (+13.4)	-	70.5 (+13.8)
MVP (Wei et al., 2022)	300	75.4 (+4.5)	-	-
<b>MILAN</b>	400	<b>79.9</b>	400	<b>84.3</b>

Table 2: Comparison of the **linear probing** top-1 accuracy on ImageNet-1K dataset. ‘‘Epochs’’ refer to the pretraining epochs of various methods. All methods adopt  $224 \times 224$  input resolution in both pretraining and linear classifier tuning.

Method	Epochs	Object detection	Instance segmentation	Semantic segmentation
		ViT-B/-L AP <sub>box</sub>	ViT-B/-L AP <sub>mask</sub>	ViT-B/-L mIoU
Supervised	-	47.9 (+4.7) / 49.3 (+6.6)	42.9 (+2.6) / 43.9 (+4.3)	47.4 (+5.3) / 49.9 (+8.0)
MoCov3	300	47.9 (+4.7) / 49.3 (+6.6)	42.7 (+2.8) / 44.0 (+4.2)	47.3 (+5.4) / 49.1 (+8.8)
DINO	300	46.8 (+5.8)	41.5 (+4.0)	47.2 (+5.5)
BEiT	300	42.6 (+10.) / 53.3 (+2.6)	38.8 (+6.7) / 47.1 (+1.1)	45.7 (+7.0) / 53.3 (+4.6)
PeCo	300	43.9 (+8.7)	39.8 (+5.7)	46.7 (+6.0)
SplitMask	300	46.8 (+5.8)	42.1 (+3.4)	45.7 (+7.0)
CAE	800	49.2 (+3.4)	43.3 (+2.2)	48.8 (+3.9)
MAE	1600	50.3 (+2.3) / 53.3 (+2.6)	44.9 (+0.6) / 47.2 (+1.0)	48.1 (+4.6) / 53.6 (+4.3)
<b>MILAN</b>	400	<b>52.6 / 55.9</b>	<b>45.5 / 48.2</b>	<b>52.7 / 57.9</b>

Table 3: Results of object detection and instance segmentation are obtained by using Mask R-CNN on COCO dataset with an input resolution of  $1024 \times 1024$ . Semantic segmentation results are obtained by using UperNet on ADE20K with an input resolution of  $512 \times 512$ . All methods use ViT models pretrained on ImageNet-1K dataset as the backbone. ‘‘Epochs’’ refer to the pretraining epochs. Previous results from He et al. (2017); Xiao et al. (2018); Chen et al. (2021); Caron et al. (2021); Bao et al. (2021); Dong et al. (2021); El-Nouby et al. (2021); Chen et al. (2022); He et al. (2021).

1600-epoch pretraining, MILAN advances AP<sub>box</sub> and AP<sub>mask</sub> by 2.3 and 0.6 points on ViT-Base but only pretrains for 400 epochs.

**Semantic segmentation on ADE20K.** We also transfer our pretrained models to semantic segmentation task on the ADE20K dataset (Zhou et al., 2017). Following the training recipe provided by MAE (He et al., 2021), the ViT models pretrained on ImageNet-1K dataset serve as the backbone of UperNet (Xiao et al., 2018), and are finetuned together with the segmentation layers. In Table 3, we report the mean intersection over union (mIoU) averaged over all semantic categories. Our method significantly improves the transferring results of ViT-Base to 52.7, surpassing MAE by 4.6 points.

### 3.3 ROBUSTNESS EVALUATION

We evaluate the robustness of our models to adversarial examples on ImageNet-Adversarial dataset (Hendrycks et al., 2021b) and distribution shifts on ImageNet-Rendition (Hendrycks et al., 2021a) and ImageNet-Sketch (Wang et al., 2019) datasets. We only finetune our pretrained models on the original ImageNet-1K training set and directly run inference on these different validation sets, without any specialized finetuning. As shown in Table 4, MILAN significantly outperforms previous state-of-the-art models. Compared with more advanced architecture RobustViT that is specially designed for robustness and is pretrained on ImageNet-22K, MILAN with vanilla ViT-Base architecture achieves accuracy gains of 7.9%~19.9% on these three datasets. When using ViT-Base, MILAN also surpasses MAE by 26.3%, 15.8% and 11.8% on these three datasets, respectively.

Method	Parameters	ImageNet-Adversarial(%)	ImageNet-Rendition(%)	ImageNet-Sketch(%)
Supervised	86M / 307M	27.2 (+35.0) / 29.6 (+46.3)	49.4 (+14.7) / 50.9 (+25.7)	35.6 (+10.7) / 37.5 (+19.9)
Swin	88M	35.8 (+26.4)	46.6 (+17.5)	32.4 (+13.9)
RobustViT	92M	42.3 (+19.9)	52.6 (+11.5)	38.4 (+7.9)
MAE	86M / 307M	35.9 (+26.3) / 57.1 (+18.8)	48.3 (+15.8) / 59.9 (+16.7)	34.5 (+11.8) / 45.3 (+12.1)
<b>MILAN</b>	86M / 307M	<b>62.2 / 75.9</b>	<b>64.1 / 76.6</b>	<b>46.3 / 57.4</b>

Table 4: Comparison of robustness to adversarial examples and distribution shifts on ImageNet. Previous results from Dosovitskiy et al. (2020); Liu et al. (2021); Mao et al. (2022); He et al. (2021).

	CLIP target	Prompting decoder	Semantic aware sampling	Epochs	Top-1 (%)
#1			Baseline (MAE)	400 (1600)	83.0 (83.6)
#2	✓			400	83.9
#3		✓		400	83.0
#4			✓	400	83.3
#5		✓	✓	400	83.3
#6	✓		✓	400	84.1
#7	✓	✓		400	85.1
#8	✓	✓	✓	400 (1600)	<b>85.4</b> (85.6)
#9	SLIP target	✓	✓	400	84.4

Table 5: Ablation study of different components in MILAN. All results are obtained by pretraining and finetuning ViT-Base model on ImageNet-1K dataset at  $224 \times 224$  resolution.

### 3.4 ABLATION STUDY

We investigate the effectiveness of the different components in MILAN through an ablation study in Table 5. More ablation studies on other tasks can be found in the appendix. Here, the results are based on pretraining a ViT-Base model on ImageNet-1K dataset for 400 epochs, followed by a 100-epoch finetuning. We tune the optimal learning rate for each entry, including the MAE baseline.

(1) By changing the reconstruction target from raw pixels to language guided representations provided by CLIP, the top-1 accuracy is improved by 0.9% (#2 vs. #1 in Table 5). We hypothesize that our target provides more semantic learning signals for pretraining, and encourages the model to get a good grasp of the visual contents instead of the low level statistics.

(2) With the CLIP target, replacing the original decoder in MAE by our prompting decoder in Figure 1(b) further improves the accuracy by 1.2% (#7 vs. #2 in Table 5). However, prompting decoder does not increase accuracy when the raw pixels are reconstructed, as the MAE model does (#3 vs. #1 in Table 5). This big difference can be explained by the different pretraining objectives. When the targets are in the latent space, the encoder’s output of the unmasked patches’ representations should be able to align with the targets without requiring further updates in the decoder. Therefore, in our case of using the CLIP target, the proposed efficient decoder brings in significant accuracy improvements. While in MAE, the encoder’s output requires transformations in the decoder to map from the latent space back to raw pixels. In short, the decoder design is heavily correlated with the reconstruction target. We realize the critical coupling effect of these two components and propose mutually beneficial design choices to boost the accuracy.

(3) The proposed semantic aware mask sampling strategy is generally beneficial regardless of the reconstruction target. After applying the CLIP target and the prompting decoder, semantic aware sampling further improves the uniformly random sampling by 0.3%, yielding the best 85.4% accuracy among different model variants (#8 vs. #7 in Table 5). Semantic aware sampling slightly reduces the pretraining difficulty and obtains lower training loss, because it favours more important image regions. Although the task becomes easier, the learned representation enjoys better accuracy.

(4) To demonstrate that masked image pretraining generally benefits from the language supervised reconstruction targets, we use a different language-image model, SLIP (Mu et al., 2021), to generate the image features as the reconstruction target in MILAN. Reconstructing the image features from SLIP still outperforms reconstructing raw pixels, surpassing MAE by 1.4% (#9 vs. #1 in Table 5). SLIP incorporates contrastive image self-supervision into language-image pretraining, and finetuning SLIP’s image encoder yields an accuracy of 82.6%. Similar to our observation on CLIP, one more step of reconstruction based pretraining by MILAN further improves the representation quality, boosting the accuracy by 1.8% compared to finetuning SLIP.

Finally, we note that a very long pretraining schedule is no longer necessary for MILAN compared to MAE. Our method enjoys much fewer epochs while achieving higher accuracy: MILAN achieves 85.4% after a 400-epoch pretraining, while MAE achieves 83.6% after a 1600-epoch pretraining.

## 4 RELATED WORKS

**Masked image pretraining.** Self-supervised pretraining aims to learn transferable representations from unlabeled data by a pretext task (Caron et al., 2018; Doersch et al., 2015; Van den Oord et al., 2018; Ermolov et al., 2021; Li et al., 2020; Goyal et al., 2021; Zbontar et al., 2021). Recent pretraining methods revitalize the use of denoising autoencoders (Vincent et al., 2010; Pathak et al., 2016; Chen et al., 2020a; Zhang et al., 2016; Noroozi & Favaro, 2016) to train the vision transformer with masked prediction objectives. The model receives incomplete images with a large portion of the patches removed and learns to reconstruct the missing contents pixel by pixel (He et al., 2021; Xie et al., 2021b; Li et al., 2021b). To inject semantics into the representations, BEiT (Bao et al., 2021), PeCo (Dong et al., 2021) and CAE (Chen et al., 2022) predict discrete visual vocabularies produced from separately trained tokenizers (Rolfe, 2016; Van Den Oord et al., 2017; Ramesh et al., 2021). MaskFeat (Wei et al., 2021) finds that the local gradient features produced by the manually-crafted HOG descriptor surpasses more complex targets. Other works like iBOT (Zhou et al., 2021), data2vec (Baevski et al., 2022), and (Kakogeorgiou et al., 2022) adopts self-distillation, where the model reconstructs the masked patches’ representations produced by the exponential moving average of the model. MILAN differs that it learns the masked autoencoder by reconstructing the latent representations that embed rich semantic information stemming from language supervision. Moreover, we find that freezing the encoder’s output features in the decoder is a critical factor when the model learns to reconstruct the latent targets. Finally, we propose a semantic aware mask sampling mechanism and alleviate the need for very long pretraining.

**Language-Image pretraining.** Learning visual representations from language supervision is not new. Early work (Frome et al., 2013) embeds images and texts into a shared semantic space so that the model is able to recognize classes even without explicit labels. Other methods leverage the caption supervision to train the vision model by completing the image captioning task (Desai & Johnson, 2021) or the masked language modeling task (Sariyildiz et al., 2020). Recently, benefiting from contrastive training (Yuan et al., 2021) and the scalability of modern backbones, CLIP (Radford et al., 2021) and ALIGN (Jia et al., 2021) learn strong visual representations on large-scale image-text datasets, advancing the transfer performance on the downstream vision tasks. Later works improve CLIP by introducing more auxiliary loss functions to assist the image-text contrastive loss, such as image self-supervision loss (Mu et al., 2021; Li et al., 2021a), self-distillation loss (Cheng et al., 2021), and token-wise max similarity (Yao et al., 2021). In this work, MILAN further improves the representation quality of language-image pretraining by incorporating masked image pretraining.

**Contrastive learning.** Contrastive methods (Caron et al., 2021; Chen et al., 2021; 2020b; He et al., 2020; Tian et al., 2020; Chen & He, 2021; Grill et al., 2020; Caron et al., 2020; Wu et al., 2018; Wang et al., 2021; Xie et al., 2021a; Alexey et al., 2016) learn augmentation invariance by enforcing similarity between different views augmented from the same image while avoiding model collapse. The learned representations show high linear separability and are commonly evaluated by linear probing. However, contrastive learning heavily depends on strong data augmentations and effective negative sampling. In contrast, MILAN uses a masked prediction objective with a reconstruction loss. Our method learns powerful representations with much simpler data augmentations.

## 5 CONCLUSION

The masked autoencoders can extract visual concepts from the unlabeled raw image pixels, which reduces the heavy reliance on large labeled datasets in computer vision tasks. However, such visual concepts may still lag the rich semantic data the image captions contain. Understanding images assisted by language captions has also been explored in the weakly supervised pretraining setting. The learned features may easily be transferred to downstream tasks via zero shot learning. However, finetuning those models directly may not reveal competitive results. In this paper, we combined these two lines of work to use the outputs of language-image pretraining as the reconstruction target for masked autoencoders, proposed a more effective decoder architecture and a semantic aware sampling mechanism. We have shown that by combining the two methods in the self-supervised pretraining, we can achieve better quality than applying each method individually.

## REFERENCES

- Dosovitskiy Alexey, Philipp Fischer, Jost Tobias, Martin Riedmiller Springenberg, and Thomas Brox. Discriminative, unsupervised feature learning with exemplar convolutional, neural networks. In *Transactions on Pattern Analysis and Machine Intelligence*, volume 38, pp. 1734–1747, 2016.
- 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. *arXiv preprint arXiv:2202.03555*, 2022.
- Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *European Conference on Computer Vision*, pp. 132–149, 2018.
- 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, pp. 9912–9924, 2020.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *International Conference on Computer Vision*, pp. 9650–9660, 2021.
- Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *International Conference on Machine Learning*, pp. 1691–1703, 2020a.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pp. 1597–1607. PMLR, 2020b.
- Xiaokang Chen, Mingyu Ding, Xiaodi Wang, Ying Xin, Shentong Mo, Yunhao Wang, Shumin Han, Ping Luo, Gang Zeng, and Jingdong Wang. Context autoencoder for self-supervised representation learning. *arXiv preprint arXiv:2202.03026*, 2022.
- Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Conference on Computer Vision and Pattern Recognition*, pp. 15750–15758, 2021.
- Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *International Conference on Computer Vision*, pp. 9640–9649, 2021.
- Ruizhe Cheng, Bichen Wu, Peizhao Zhang, Peter Vajda, and Joseph E Gonzalez. Data-efficient language-supervised zero-shot learning with self-distillation. In *Conference on Computer Vision and Pattern Recognition*, pp. 3119–3124, 2021.
- Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition*, pp. 248–255. Ieee, 2009.
- Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. In *Conference on Computer Vision and Pattern Recognition*, pp. 11162–11173, 2021.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *International Conference on Computer Vision*, pp. 1422–1430, 2015.

- Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. Peco: Perceptual codebook for bert pre-training of vision transformers. *arXiv preprint arXiv:2111.12710*, 2021.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jegou, and Edouard Grave. Are large-scale datasets necessary for self-supervised pre-training? *arXiv preprint arXiv:2112.10740*, 2021.
- Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. In *International Conference on Machine Learning*, pp. 3015–3024, 2021.
- Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems*, volume 26, 2013.
- Priya Goyal, Mathilde Caron, Benjamin Lefaudeaux, Min Xu, Pengchao Wang, Vivek Pai, Mannat Singh, Vitaliy Liptchinsky, Ishan Misra, Armand Joulin, et al. Self-supervised pretraining of visual features in the wild. *arXiv preprint arXiv:2103.01988*, 2021.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 21271–21284, 2020.
- Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *International Conference on Computer Vision*, pp. 2961–2969, 2017.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8340–8349, 2021a.
- Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15262–15271, 2021b.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pp. 4904–4916. PMLR, 2021.
- Ioannis Kakogeorgiou, Spyros Gidaris, Bill Psomas, Yannis Avrithis, Andrei Bursuc, Konstantinos Karantzas, and Nikos Komodakis. What to hide from your students: Attention-guided masked image modeling. *arXiv preprint arXiv:2203.12719*, 2022.
- Junnan Li, Pan Zhou, Caiming Xiong, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*, 2020.
- Xiaotong Li, Yixiao Ge, Kun Yi, Zixuan Hu, Ying Shan, and Ling-Yu Duan. mc-beit: Multi-choice discretization for image bert pre-training. *arXiv preprint arXiv:2203.15371*, 2022.

- Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021a.
- Zhaowen Li, Zhiyang Chen, Fan Yang, Wei Li, Yousong Zhu, Chaoyang Zhao, Rui Deng, Liwei Wu, Rui Zhao, Ming Tang, et al. Mst: Masked self-supervised transformer for visual representation. In *Advances in Neural Information Processing Systems*, volume 34, 2021b.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pp. 740–755. Springer, 2014.
- Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Conference on Computer Vision and Pattern Recognition*, pp. 2117–2125, 2017.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10012–10022, 2021.
- Xiaofeng Mao, Gege Qi, Yuefeng Chen, Xiaodan Li, Ranjie Duan, Shaokai Ye, Yuan He, and Hui Xue. Towards robust vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12042–12051, 2022.
- Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. *arXiv preprint arXiv:2112.12750*, 2021.
- Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pp. 69–84. Springer, 2016.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Conference on Computer Vision and Pattern Recognition*, pp. 2536–2544, 2016.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831. PMLR, 2021.
- Jason Tyler Rolfe. Discrete variational autoencoders. *arXiv preprint arXiv:1609.02200*, 2016.
- Mert Bulent Sariyildiz, Julien Perez, and Diane Larlus. Learning visual representations with caption annotations. In *European Conference on Computer Vision*, pp. 153–170. Springer, 2020.
- Mannat Singh, Laura Gustafson, Aaron Adcock, Vinicius de Freitas Reis, Bugra Gedik, Raj Prateek Kosaraju, Dhruv Mahajan, Ross Girshick, Piotr Dollár, and Laurens van der Maaten. Revisiting weakly supervised pre-training of visual perception models. *arXiv preprint arXiv:2201.08371*, 2022.
- Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *International Conference on Computer Vision*, pp. 843–852, 2017.
- Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? In *Advances in Neural Information Processing Systems*, volume 33, pp. 6827–6839, 2020.
- Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit. *arXiv preprint arXiv:2204.07118*, 2022.

- Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Advances in neural information processing systems*, volume 30, 2017.
- Aaron Van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. In *Journal of Machine Learning Research*, volume 11, 2010.
- Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. *Advances in Neural Information Processing Systems*, 32, 2019.
- Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Conference on Computer Vision and Pattern Recognition*, pp. 3024–3033, 2021.
- Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. *arXiv preprint arXiv:2112.09133*, 2021.
- Longhui Wei, Lingxi Xie, Wengang Zhou, Houqiang Li, and Qi Tian. Mvp: Multimodality-guided visual pre-training. *arXiv preprint arXiv:2203.05175*, 2022.
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Conference on Computer Vision and Pattern Recognition*, pp. 3733–3742, 2018.
- Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *European Conference on Computer Vision*, pp. 418–434, 2018.
- Zhenda Xie, Yutong Lin, Zheng Zhang, Yue Cao, Stephen Lin, and Han Hu. Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In *Conference on Computer Vision and Pattern Recognition*, pp. 16684–16693, 2021a.
- 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. *arXiv preprint arXiv:2111.09886*, 2021b.
- Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021.
- Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.
- Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- Xin Yuan, Zhe Lin, Jason Kuen, Jianming Zhang, Yilin Wang, Michael Maire, Ajinkya Kale, and Baldo Faieta. Multimodal contrastive training for visual representation learning. In *Conference on Computer Vision and Pattern Recognition*, pp. 6995–7004, 2021.
- Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pp. 12310–12320, 2021.
- Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European Conference on Computer Vision*, pp. 649–666. Springer, 2016.
- Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Conference on Computer Vision and Pattern Recognition*, 2017.
- Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021.

## A APPENDIX

### A.1 IMPLEMENTATION DETAILS

**Pretraining on ImageNet-1K.** We follow the pretraining recipe of MAE (He et al., 2021), using a publicly released codebase<sup>1</sup>. We only pretrain ViT models for 400 epochs. We use AdamW optimizer with a momentum of  $(\beta_1 = 0.9, \beta_2 = 0.95)$ , a mini-batch size of 4096 and an initial learning rate of  $2.4e - 3$  (scaled based on  $lr = \text{base\_lr} \times \text{batchsize}/256$  where  $\text{base\_lr} = 1.5e - 4$ ). The learning rate is linearly warmed up for the first 40 epochs, and decayed to zero by a cosine learning rate schedule. The weight decay is set to 0.05. For data augmentation, we only adopt random resized crop to  $224 \times 224$  resolution, random horizontal flip, and normalization. The masking ratio is set to 75%. We use the image features generated from pretrained CLIP models with ViT image encoders as the reconstruction targets. The wall-clock pretraining time of MILAN for ViT-Base and ViT-Large models take about 2 days and 3 days on a machine with 8 A100 GPUs, respectively. For reference, MAE requires 1600-epoch pretraining which takes 116 hours on ViT-Base model.

**Finetuning on ImageNet-1K.** We follow the finetuning recipe of MAE (He et al., 2021) but tune the optimal learning rates for our models. We only finetune 100 epochs. We use AdamW optimizer with a momentum of  $(\beta_1 = 0.9, \beta_2 = 0.999)$ , a mini-batch size of 1024 and an initial learning rate of  $4e - 4$  (scaled based on  $lr = \text{base\_lr} \times \text{batchsize}/256$  where  $\text{base\_lr} = 1e - 4$ ). The learning rate is linearly warmed up for the first 5 epochs, and decayed to zero by a cosine learning rate schedule. The layer-wise learning rate decay factor is set to 0.65 for ViT-Base and 0.75 for ViT-Large. The weight decay is set to 0.05. We adopt RandAugment (Cubuk et al., 2020), and set  $\text{label\_smoothing} = 0.1$ ,  $\text{mixup} = 0.8$ ,  $\text{cutmix} = 1.0$ ,  $\text{drop\_path} = 0.1$  (ViT-Base),  $0.2$  (ViT-Large).

**Linear probing on ImageNet-1K.** We follow the linear probing recipe of MAE (He et al., 2021). We train the linear classifier for 100 epochs. We use LARS (You et al., 2017) optimizer with a momentum of 0.9, a mini-batch size of 16384, and an initial learning rate of 3.2 for ViT-Base and 1.28 for ViT-Large. The learning rate is linearly warmed up for the first 10 epochs, and decayed to zero by a cosine learning rate schedule. We do not use mixup, cutmix, drop path, or color jittering, and the weight decay is set to zero.

**Objection detection and instance segmentation on COCO.** Following (He et al., 2021), the pretrained ViT models by MILAN are adapted to FPN (Lin et al., 2017) in the Mask R-CNN framework (He et al., 2017), and we finetuned end-to-end on COCO training set (Lin et al., 2014). We use AdamW optimizer with a momentum of  $(\beta_1 = 0.9, \beta_2 = 0.999)$ , a mini-batch size of 16, and an initial learning rate of  $2e - 4$ . The layer-wise learning rate decay is set to 0.75 for ViT-Base and 0.85 for ViT-Large. The weight decay is set to 0.1 and the drop path rate is set to 0.1 for ViT-Base and 0.2 for ViT-Large. We adopt the standard  $1 \times$  schedule: 12 epochs with the learning rate decayed by 10 at epochs 8 and 11. The input resolution is  $1024 \times 1024$ . We do not use multi-scale testing. We build upon a public codebase<sup>2</sup> that reproduces MAE’s detection results.

**Semantic segmentation on ADE20K.** Following (He et al., 2021), the ViT models pretrained on ImageNet-1K dataset by MILAN serve as the backbone of UperNet framework (Xiao et al., 2018), and are finetuned together with the segmentation layers on ADE20K dataset (Zhou et al., 2017) for 160K iterations. We use AdamW optimizer with a momentum of  $(\beta_1 = 0.9, \beta_2 = 0.999)$ , a mini-batch size of 16 and an initial learning rate of  $3e - 5$ . The learning rate is linearly warmed up for the first 1500 iterations, and decayed to zero by a poly learning rate schedule. The layer-wise learning rate decay is set to 0.9. The weight decay is set to 0.05 and the drop path rate is set to 0.1. The input resolution is  $512 \times 512$ . We do not use multi-scale testing. We build upon a public codebase<sup>3</sup> that reproduces MAE’s segmentation results.

### A.2 MORE RESULTS

**Ablation study on semantic segmentation task.** We also conduct an ablation study the different components of MILAN on the semantic segmentation task, as shown in Table 6. The results are based on pretraining a ViT-Base model on ImageNet-1K for 400 epochs, followed by finetuning UperNet

<sup>1</sup><https://github.com/facebookresearch/mae>

<sup>2</sup><https://github.com/hustvl/MIMDet>

<sup>3</sup>[https://github.com/implus/mae\\_segmentation](https://github.com/implus/mae_segmentation)

	Method	mIoU
#1	Baseline (MAE)	48.1
#2	+ CLIP target	49.2
#3	+ CLIP target + Semantic aware sampling	50.4
#4	+ CLIP target + Prompting decoder	51.9
#5	+ CLIP target + Prompting decoder + Semantic aware sampling	<b>52.7</b>

Table 6: Ablation study on different components of MILAN on the semantic segmentation task. All results are based on ViT-Base models that are pretrained on ImageNet-1K and finetuned on the ADE20K dataset using the UperNet segmentation framework.

	Method	Top-1 (%)
#1	Baseline (MAE)	62.0
#2	+ CLIP target	67.1
#3	+ CLIP target + Semantic aware sampling	68.1
#4	+ CLIP target + Prompting decoder	79.9
#5	+ CLIP target + Prompting decoder + Semantic aware sampling	78.9
#6	#5 + Semantic aware probing	<b>80.0</b>

Table 7: Ablation study on different components of MILAN on the linear probing task. All results are based on ViT-Base models that are pretrained on ImageNet-1K dataset at  $224 \times 224$  resolution.

on ADE20K by 160K iterations. We find that the overall trend is consistent with our findings in finetuning results on ImageNet. By changing the reconstruction target from raw pixels to image features produced from CLIP, the mIoU is improved by 1.1 points (#2 vs. #1 in Table 6). With the CLIP target, replacing the random masking in MAE by our semantic aware sampling (#3 vs. #2 in Table 6) or replacing the vanilla decoder in MAE by our prompting decoder (#4 vs. #2 in Table 6) further improves the mIoU by 1.2 and 2.7 points, respectively. Finally, applying both the semantic aware sampling and the prompting decoder leads to the best 52.7 mIoU, which is 3.5 points higher than applying the CLIP target alone and 4.6 points higher than the MAE baseline.

**Ablation study on linear probing task.** We further conduct an ablation study on the different components of MILAN on linear probing task, as shown in Table 7. The results are based on pretraining a ViT-Base model on ImageNet-1K dataset for 400 epochs, followed by a 100-epoch linear classifier training. Consistent with our findings on finetuning and semantic segmentation tasks, using CLIP target brings in accuracy improvement. By changing the reconstruction target from raw pixels to image features produced from CLIP with language supervision, the top-1 accuracy is boosted by 5% (#2 vs. #1 in Table 7). With the CLIP target, replacing the random masking in MAE by our semantic aware mask sampling (#3 vs. #2 in Table 7) or replacing the original decoder in MAE by our prompting decoder (#4 vs. #2 in Table 7) improves the accuracy by 1% and 12.8%, respectively. However, inconsistent with our findings on finetuning and semantic segmentation tasks, applying both the prompting decoder and the semantic aware mask sampling simultaneously does not lead to the best accuracy (#5 vs. #4 in Table 7). We hypothesize that the encoder model is overly adapted to those unmasked important image patches when semantic aware sampling and prompting decoder are applied together. In linear probing, the pretrained encoder model receives full patches, and its weights are frozen, with only the linear classifier’s weights being updated. Thus, the model may not be able to cope with the image patches that are less relevant or irrelevant to the objects that need to be classified, and those features degrade the performance of the linear classifier.

To support our speculation that the pretrained encoder may be biased towards important patches, we perform a semantic aware probing experiment, listed in #6 of Table 7. Specifically, the pretrained model is obtained by MILAN with CLIP target, prompting decoder and semantic aware sampling. When finetuning the linear classifier as well as performing inference on the validation dataset, we select the top 50% important image patches. Only the selected patches are fed into the frozen encoder model to obtain the features to train the linear classifier. Although the classifier is trained on features from incomplete inputs, semantic aware probing improves the accuracy to 80% (#6 vs. #5 in Table 7),

Method	ViT-Base		ViT-Large	
	Epochs	Top-1 (%)	Epochs	Top-1 (%)
<i>contrastive or clustering based</i>				
MoCov3 (Chen et al., 2021)	300	76.7 (+3.2)	300	77.6 (+6.7)
DINO (Caron et al., 2021)	400	78.2 (+1.7)	-	-
iBoT (Zhou et al., 2021)	1600	79.5 (+0.4)	1000	81.0 (+3.3)
<i>reconstruction based</i>				
BEiT (Bao et al., 2021)	800	56.7 (+23.2)	800	73.5 (+10.8)
SimMIM (Xie et al., 2021b)	800	56.7 (+23.2)	-	-
MaskFeat (Wei et al., 2021)	-	-	1600	67.7 (+16.6)
CAE (Chen et al., 2022)	800	68.3 (+11.6)	-	-
MAE (He et al., 2021)	1600	68.0 (+11.9)	1600	75.8 (+8.5)
<i>language-image pretraining based</i>				
CLIP (Radford et al., 2021)	-	66.5 (+13.4)	-	70.5 (+13.8)
MVP (Wei et al., 2022)	300	75.4 (+4.5)	-	-
<b>MILAN w/ SAS</b>	400	78.9 (+1.0)	400	84.1 (+0.2)
<b>MILAN w/ RS</b>	400	<b>79.9</b>	400	<b>84.3</b>

Table 8: Comparison of the linear probing top-1 accuracy on ImageNet-1K dataset. ‘‘SAS’’: semantic aware sampling. ‘‘RS’’: random sampling. ‘‘Epochs’’ refer to the pretraining epochs of various methods. All methods adopt  $224 \times 224$  input resolution in pretraining and linear probing.

Method	Epochs	Top-1 (%)
KD	400	84.0
MILAN	400	85.4

Table 9: Compare MILAN with knowledge distillation (KD) on ImageNet-1K using ViT-B/16 model. Both methods use the representations produced by the CLIP image encoder as the target features. We report the finetuning top-1 accuracy. ‘‘Epoch’’ refers to the pretraining epochs.

indicating that the pretrained encoder model is more adept at extracting features from semantically important patches.

From our ablation study on the linear probing task, we find that MILAN with random masking gives better accuracy the ViT-Base model. In Table 8, we include linear probing results obtained by applying MILAN with random masking on both ViT-Base and ViT-Large models. Compared with MILAN with semantic aware sampling, the linear probing accuracies are further improved by 1% and 0.2%, respectively. Compared with the state-of-the-art contrastive method (Zhou et al., 2021), MILAN with random sampling achieves higher linear probing accuracy on both ViT-Base (+0.4%) and ViT-Large (+3.3%) models.

**MILAN vs. knowledge distillation.** In MILAN, the decoder reconstructs the representations of the masked patches with the assistance from the encoder’s output features of the unmasked patches. The reconstruction loss is computed on both the encoder’s output features of the unmasked patches and decoder’s output features of the masked patches. Here, we perform another ablation by removing the decoder from MILAN. The pretraining objective becomes training the encoder only to predict the target features on the unmasked patches, which can be regarded as a semantic aware knowledge distillation (KD) method. For KD, we also use the image features from the pretrained CLIP image encoder as the target, and the pretraining loss is the mean squared error between the normalized target features and the encoder’s output features. After pretraining, the encoder model is finetuned end-to-end with the same recipe a MILAN. The results are shown in Table 9. KD achieves 84.0% top-1 accuracy, while direct finetuning of the CLIP image encoder only yields 82.1% accuracy. Due to the possible data distribution gap between the pretraining data (OpenAI’s 400M) and the finetuning data (ImageNet-1K), finetuning the CLIP image encoder may not lead to a strong performance. But the gap can be overcome by pretraining the model with a KD objective on ImageNet-1K followed by finetuning. Moreover, MILAN achieves 1.4% higher accuracy than KD. MILAN creates a more challenging pretraining task, where the model not only needs to predict the latent features of the

CLIP image encoder	FT on IN1K top-1 (%)	LP on IN1K top-1 (%)	OD on COCO AP <sub>box</sub>	IS on COCO AP <sub>mask</sub>	SS on ADE20K mIoU
ViT-Base	86.7	81.6	55.0	47.5	55.3
ViT-Large	87.8	84.3	55.9	48.2	57.9

Table 10: Comparison of using different CLIP image encoders to produce the reconstruction targets for pretraining ViT-Large. We compare the results of finetuning (FT) and linear probing (LP) on ImageNet-1K (IN1K), object detection (OD) and instance segmentation (IS) on COCO, and semantic segmentation (SS) on ADE20K.

visible unmasked patches but also learns to reconstruct the representations of the invisible masked patches through the decoder. The masked image reconstruction task can better leverage the guidance from the target features and improve the representation quality.

**Impact of different CLIP image encoders.** In Table 10, we compare the results of using ViT-Base version and ViT-Large version of the CLIP image encoder to pretrain the ViT-Large model using our MILAN method. The results are obtained by 400 epochs of pretraining on ImageNet-1K, followed by finetuning, linear probing or transfer learning, using the same procedures as described in Appendix A.1. Using image features produced from the ViT-L CLIP image encoder as the targets consistently improves the performance on all tasks. For example, it achieves 1.1% higher accuracy than using the ViT-Base CLIP image encoder on ImageNet finetuning. The number 87.8% improves our ViT-Base result in Table 1 by 2.4% on ImageNet finetuning, showing that MILAN scales well with model size.

### A.3 VISUALIZATIONS

In Figure 3, we provide visualizations of the learned representations from MAE, MAE with CLIP reconstruction target, and our MILAN method. MILAN can better extract the important visual contents inside the images compared to both MAE and MAE with CLIP target, indicating that the proposed prompting decoder and semantic aware sampling contribute significantly to learning higher quality visual representations on top of the CLIP target. Moreover, in Figure 4, we show that the proposed semantic aware sampling indeed favours more important image regions. The 25% unmasked patches cover the contents that are more related to the objects in the images. The semantic aware sampling facilitates the model to learn better on more important foreground regions, leading to accuracy improvements on finetuning, linear probing and semantic segmentation tasks as shown in our ablation studies (see Tables 5,6,7).

### A.4 LIMITATION

Similar to (Bao et al., 2021; Chen et al., 2022; Li et al., 2022; Wei et al., 2022) which rely on external datasets to train their image tokenizers, the reconstruction target in MILAN is obtained from the CLIP model which also requires an extra image-text dataset. Training the CLIP model, if it is not amortized for many downstream tasks, is considered an extra training step. In practice, we use publicly available pretrained CLIP models, so our method does not require a bespoke CLIP training step. Moreover, we only perform inference on the CLIP image encoders to produce the reconstruction targets, which are only used in the pretraining phase. CLIP is not used in finetuning or linear probing stages, regardless of the classification, detection or segmentation tasks. The language-image (LM) models like CLIP are becoming important and popular pretrained models to be applied to downstream tasks. Our method can be regarded as a useful intermediate step, as our obtained models notably outperform these strong LM models on various tasks.

### A.5 SOCIETAL IMPACTS

The proposed MILAN method produces transferable representations based on the learned statistics of the training dataset. Therefore, the trained model may also reflect the biases in the training data. Moreover, since MILAN uses the image features generated from the CLIP model as the reconstruction targets and the CLIP model itself is trained on an uncurated image-text dataset containing English-only captions, performance on images collected from non-English speaking countries requires further

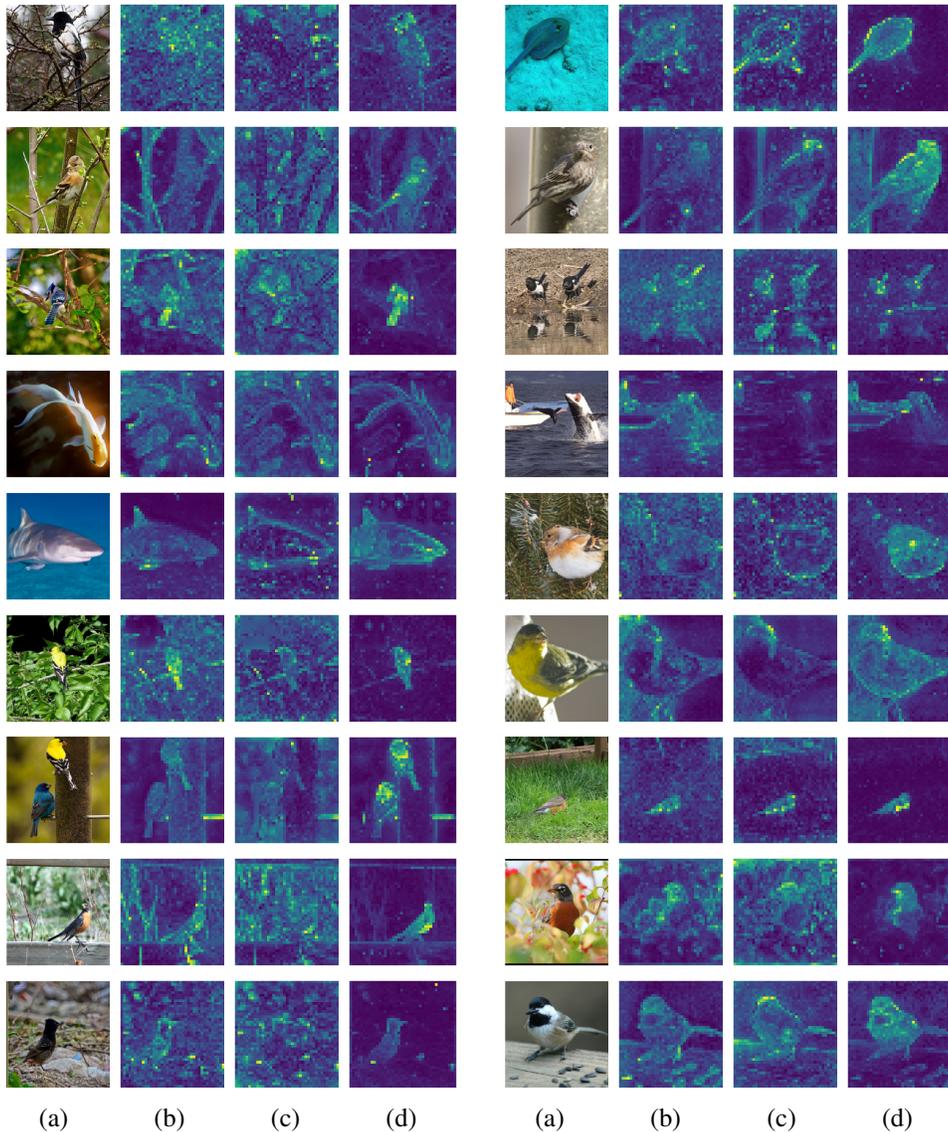


Figure 3: Visualizations of (a) original images, (b) the attention features extracted from the last self-attention layer of the ViT-Base model pretrained by MAE, (c) MAE with CLIP reconstruction target, and (d) our MILAN method. MILAN can better extract the important visual contents inside the images compared to both MAE and MAE with CLIP target.



Figure 4: Visualization of the original images (left), masked images by the semantic aware sampling strategy with 75% masking ratio (middle), and the reconstruction loss patch-by-patch (right). For the plots of reconstruction loss, darker green colors indicate higher loss values. As shown, both unmasked patches and masked foreground patches have lower losses.

research. In future works, one may apply the MILAN method by taking multi-lingual language assisted representations as the reconstruction target. Extension to large-scale transformer model pretraining on video datasets using the MILAN framework could also be a future direction.