

Semantic-Aware Autoregressive Image Modeling for Visual Representation Learning

Kaiyou Song*, Shan Zhang, Tong Wang

Megvii Technology
{songkaiyou, zhangshan, wangtong}@megvii.com

Abstract

The development of autoregressive modeling (AM) in computer vision lags behind natural language processing (NLP) in self-supervised pre-training. This is mainly caused by the challenge that images are not sequential signals and lack a natural order when applying autoregressive modeling. In this study, inspired by human beings' way of grasping an image, i.e., focusing on the main object first, we present a semantic-aware autoregressive image modeling (SemAIM) method to tackle this challenge. The key insight of SemAIM is to autoregressive model images from the semantic patches to the less semantic patches. To this end, we first calculate a semantic-aware permutation of patches according to their feature similarities and then perform the autoregression procedure based on the permutation. In addition, considering that the raw pixels of patches are low-level signals and are not ideal prediction targets for learning high-level semantic representation, we also explore utilizing the patch features as the prediction targets. Extensive experiments are conducted on a broad range of downstream tasks, including image classification, object detection, and instance/semantic segmentation, to evaluate the performance of SemAIM. The results demonstrate SemAIM achieves state-of-the-art performance compared with other self-supervised methods. Specifically, with ViT-B, SemAIM achieves 84.1% top-1 accuracy for fine-tuning on ImageNet, 51.3% AP and 45.4% AP for object detection and instance segmentation on COCO, which outperforms the vanilla MAE by 0.5%, 1.0%, and 0.5%, respectively. Code is available at <https://github.com/skyoux/SemAIM>.

Introduction

With the rapid development of masked language modeling (MLM) (e.g., BERT (Devlin et al. 2018)) and autoregressive language modeling (ALM) (e.g., GPT (Radford et al. 2018, 2019; Brown et al. 2020)), self-supervised pre-training has achieved impressive performance in learning extensible representations in the field of natural language processing (NLP). Recently, inspired by the masking mechanism in MLM, masked image modeling (MIM) (Bao, Dong, and Wei 2022; He et al. 2022; Zhou et al. 2022; Xie et al. 2022) has been proposed and rapidly improved in the computer vision community. The key to applying MIM is to use a high mask

ratio on images to reduce spatial redundancy. MIM achieves a better performance of self-supervised learning (SSL) in many downstream tasks (Russakovsky et al. 2015; Lin et al. 2014; Zhou et al. 2017) compared with other alternatives (He et al. 2020; Chen et al. 2020b; Chen, Xie, and He 2021; Caron et al. 2021). By contrast, the development of autoregressive image modeling (AIM) lags behind MIM in computer vision due to the significant difference between language and vision.

As discussed in (He et al. 2022), languages are human-generated signals with high-level semantics and dense information. They are sequential signals and provide a natural order for applying autoregressive modeling. For example, the "left-to-right" ALM used in GPT (Radford et al. 2018, 2019; Brown et al. 2020) shows a strong language generation and understanding ability. On the contrary, images are natural signals with heavy spatial redundancy. They are not sequential signals and lack a natural order for applying autoregressive modeling. Therefore, autoregressive image modeling is not an effortless way compared to masked image modeling. Several studies (Chen et al. 2020a; Hua et al. 2022; Qi et al. 2022) attempt to conduct autoregressive image modeling for self-supervised learning. They propose to use raster order image pixels (Chen et al. 2020a) and stochastic order image patches (Hua et al. 2022; Qi et al. 2022) for autoregressive modeling. However, we argue that raster and stochastic orders are not ideal for visual representation learning since they are inconsistent with the human visual mechanisms of grasping an image. According to the attention mechanism of the human visual system, human beings always selectively attend to the most informative parts of visual stimuli (Eriksen and Hoffman 1972; Koch et al. 2006). Specifically, they first focus on the main object or the object they are interested in, then focus on other contents in images, such as the background and other objects. This attention mechanism makes human beings grasp an image quickly.

Can we apply autoregressive image modeling by mimicking human visual mechanisms of grasping an image (i.e., focusing on the main object first). In this study, we design a semantic-aware autoregressive image modeling (SemAIM) method to answer this question. The core concept of SemAIM is to model images from the semantic patches to the less semantic patches autoregressively. To achieve this goal, we first calculate a semantic-aware order of patches according to their feature similarities and then perform

*Corresponding author (songkaiyou@foxmail.com).
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the autoregression procedure based on the permutation. Furthermore, considering that the raw pixels of patches are low-level signals and are not ideal prediction targets for learning high-level semantic representation, we also explore utilizing the patch features as the prediction targets. We conducted extensive experiments on a broad range of downstream tasks, including image classification, object detection, and instance/semantic segmentation, to evaluate the performance of SemAIM. The experimental results show SemAIM achieves state-of-the-art performance compared with other self-supervised methods, especially in dense prediction tasks. Specifically, with ViT-B, SemAIM achieves 84.1% top-1 accuracy for fine-tuning on ImageNet, 51.3% AP on COCO for object detection, and 45.4% AP on COCO for instance segmentation, which outperforms the vanilla MAE by 0.5%, 1.0%, and 0.5%, respectively. This study empirically demonstrates that SemAIM is more appropriate for autoregressive image modeling and is helpful for learning semantic visual representation.

Related Work

Self-Supervised Learning. Self-supervised learning aims to learn scalable visual representations without any human annotations. The key to self-supervised learning is how to design an effective pretext task to learn semantic representations. In the field of computer vision, early studies designed various pretext tasks, including image inpainting (Pathak et al. 2016), colorization (Zhang, Isola, and Efros 2016), jigsaw puzzle (Noroozi and Favaro 2016), counting (Noroozi, Pirsavash, and Favaro 2017), and rotation prediction (Komodakis and Gidaris 2018). Though with inferior performance, these studies laid the foundation for the development of this field. After that, contrastive learning (He et al. 2020; Chen et al. 2020b; Caron et al. 2020; Chen and He 2021; Chen, Xie, and He 2021; Caron et al. 2021; Guo et al. 2022; Song et al. 2023a,b), as a type of instance discriminative (Wu et al. 2018) method, is heavily studied and has shown remarkable progress in recent years, which aims at pulling different augmented versions of the same image closer while pushing diverse images far from each other. However, contrastive learning needs curated data and carefully-designed data augmentation techniques for pre-training (Chen et al. 2020b; Caron et al. 2021).

Masked Image Modeling. Inspired by mask language modeling (Devlin et al. 2018) in NLP, masked image modeling (Bao, Dong, and Wei 2022; He et al. 2022) has been proposed for visual pre-training in the recent two years. Many studies are proposed, which mainly focus on improving the performance of masked image modeling from two aspects, i.e., prediction targets and masking strategy. For prediction targets, various contents are explored, including raw RGB pixels (He et al. 2022; Xie et al. 2022; Chen et al. 2022; Wang et al. 2023a), discrete tokens (Bao, Dong, and Wei 2022; Dong et al. 2023), HoG features (Wei et al. 2022), features extracted from a momentum model (Zhou et al. 2022; Dong et al. 2022), and features extracted from pre-trained models (Li et al. 2022a; Hou et al. 2023). For the masking strategy, random masking (He et al. 2022; Xie et al. 2022) and block-wise masking strategy (Bao, Dong, and Wei 2022;

Zhou et al. 2022) are adopted in many methods. While some studies (Kakogeorgiou et al. 2022; Li et al. 2022a; Wang et al. 2023a) explore adaptive or learnable masking strategies to force the model to learn semantic visual representations.

Autoregressive Modeling. Autoregressive language modeling (ALM) proposed in GPT (Radford et al. 2018, 2019; Brown et al. 2020) has shown significant success in learning general representations in the field of NLP. ALM in GPT predicts the next possible word based on the preceding word in left-to-right order. Besides left-to-right prediction, permuted ALM (Yang et al. 2019; Song et al. 2020) aims to learn contextual information by maximizing the expected logarithmic likelihood of all possible permutations of sequences. While UniLM (Dong et al. 2019) adopts both left-to-right and right-to-left predictions to boost the performance of pre-trained language models.

While in the field of computer vision, autoregressive image modeling (AIM) is employed for several tasks, such as image generation (Van den Oord et al. 2016; Ramesh et al. 2021), object detection (Chen et al. 2021), and representation learning (Oord, Li, and Vinyals 2018; Chen et al. 2020a; Hua et al. 2022; Qi et al. 2022). For the representation learning task, early study CPC (Oord, Li, and Vinyals 2018) uses autoregressive models to learn representations by predicting the future in latent space. iGPT (Chen et al. 2020a) directly applies GPT (Radford et al. 2018) on images by autoregressive modeling raw pixels in raster order. Serializing images into pixels is not an ideal strategy, which greatly limits the efficiency and representation learning performance of iGPT. After that, with the development of vision transformer (Vaswani et al. 2017; Dosovitskiy et al. 2020), which serializes images into patches, recent studies (Hua et al. 2022; Qi et al. 2022) conduct autoregressive modeling based on image patches. RandSAC (Hua et al. 2022) groups patches into hierarchically arranged segments and performs autoregressive prediction on these segments in stochastic order. SAIM (Qi et al. 2022) directly applies autoregressive prediction on patches in stochastic order. However, we argue that raster or stochastic orders are not ideal for visual representation learning since they are inconsistent with the human visual mechanisms of grasping an image, i.e., focusing on the main object first. In this study, we present the semantic-aware autoregressive image modeling (SemAIM) to tackle this problem.

Method

Preliminary: Autoregressive Modeling

Autoregressive modeling aims to learn a good visual representation from an unlabeled dataset by modeling its distribution. Specifically, given an unlabeled dataset \mathcal{D} consisting of high dimensional data $\mathbf{x} = [x_1, x_2, \dots, x_N]$, a permutation π of the set $[1, N]$ can be picked, and π_i and $\pi_{<i}$ denote the i -th element and the first $1-i$ elements of the permutation. Autoregressive modeling learns the data distribution by maximizing the likelihood function:

$$\mathcal{L} = - \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \sum_{i=1}^N \log p_{\theta}(x_{\pi_i} | x_{\pi_{<i}}) \quad (1)$$

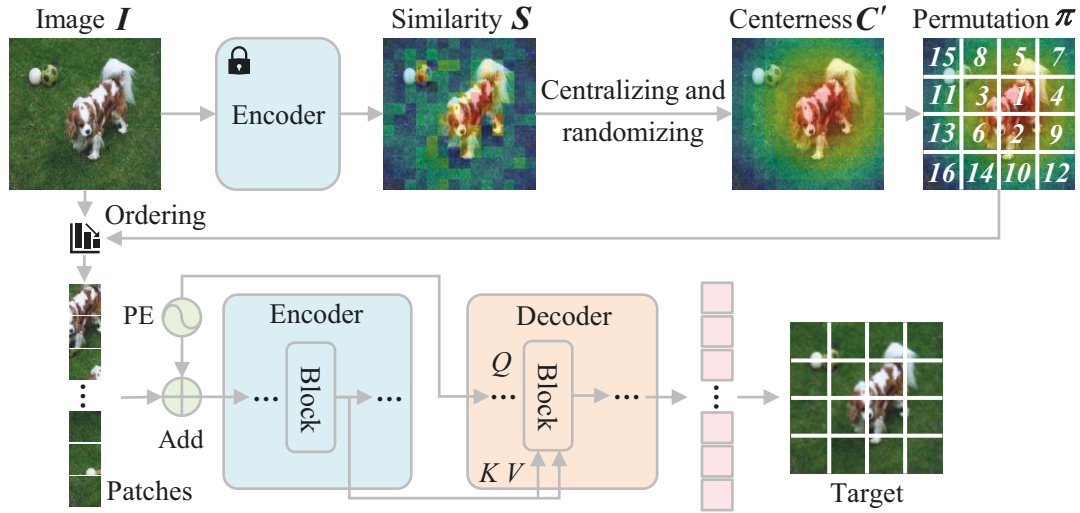


Figure 1: Illustration of SemAIM. Given an input image I , we first calculate its similarity map S and generate a semantic-aware permutation π . Then, we employ a parallel encoder-decoder for autoregressive modeling according to the permutation. “PE” denotes the position embedding, and “add” denotes element-wise addition. Note that the centerness in Eq. (5) is processed by $C' = 1 - \text{softmax}(C)$ for visualization. At fine-tuning stage, the encoder is applied for downstream tasks.

where θ is the parameters of the autoregressive model. When working with images, an image is reshaped into a sequence of pixels (Chen et al. 2020a) or patches (Hua et al. 2022; Qi et al. 2022), and the permutation is generated by a fixed raster order (Chen et al. 2020a) or a stochastic order (Hua et al. 2022; Qi et al. 2022). However, as we have analyzed before, such orders are inconsistent with the human visual understanding, i.e., focusing on the semantic object first.

Semantic-guided Autoregressive Image Modeling

In this study, we present the semantic-aware autoregressive image modeling (SemAIM) to overcome the limitations of existing autoregressive image modeling methods. Fig. 1 illustrates our proposed SemAIM. In SemAIM, we first calculate its similarity map first and generate semantic-aware permutation of its patches. Then, we employ a parallel encoder-decoder for autoregressive modeling according to the generated permutation.

Given an image $I \in \mathbb{R}^{H \times W \times C}$, it is first reshaped into a sequence of patches $I_p \in \mathbb{R}^{N \times (P^2 C)}$, where (H, W) indicates the spatial resolution, C is the number of channels, P is the patch size, and $N = HW/P^2$ is the number of patches. A linear projection is then applied to I_p , mapping it to D dimensions to get patch embeddings $x_p \in \mathbb{R}^{N \times D}$. A [CLS] token $x_{cls} \in \mathbb{R}^D$ is used to aggregate the information. 2D sin-cos position embeddings $p \in \mathbb{R}^{(N+1) \times D}$ are added to the patch embeddings to retain positional information. Thus, the initialized sequence $x = [x_{cls}; x_p] \oplus p$ can be obtained. Where \oplus denotes element-wise addition.

Semantic-aware Permutation Generation To generate semantic-aware order, we need to locate semantic regions first. In this study, we found that the similarities among the [CLS] token and patch tokens from the deep layers

of the pre-trained encoder can locate the semantic regions of input images. Therefore, we first generate semantic-aware permutation according to the similarity map.

Feed the embedded tokens x into the frozen encoder, we can get the output tokens $z = [z_{cls}, z_p]$ from the last blocks of the encoder. Then, the similarities among the [CLS] token z_{cls} and patch tokens z can be calculated:

$$S_i = \frac{\exp(\cos(z_{cls}, z_i))}{\sum_{j=1}^N \exp(\cos(z_{cls}, z_j))} \quad (2)$$

where $S \in \mathbb{R}^N$ denotes the similarity matrix, and $\cos()$ denotes cosine similarity between two vectors. Concretely, S_i means the similarity between the [CLS] token and the i -th patch token. We can reshape S into two-dimension map $S \in \mathbb{R}^{H' \times W'}$, where $H' = H/P$, $W' = W/P$. The similarity map S highlights the semantic regions in the image since the [CLS] token aggregates the global information.

In this study, we found that it is not ideal to use the similarity map directly for autoregression permutation due to two reasons. First, the similarity map is not an accurate semantic segmentation map, which is noisy and not a smooth permutation. Second, using the similarity map directly will significantly decrease the diversity of autoregression. Therefore, based on the similarity map, we further design a center-to-outward permutation motivated by human visual mechanisms of grasping an image. We can get the semantic region center according to the similarity map S :

$$c_y, c_x = \arg \max (S_{ij}) \quad (3)$$

Note that we adopt a 3×3 mean filtering operation on S to alleviate the influence of noise. The patch with index c_y, c_x is termed as the center patch. Then, the distances between patches and center are calculated:

$$D_{ij} = \sqrt{(c_y - i)^2 + (c_x - j)^2} \quad (4)$$

Not that patches located on the same radius have the same distance and show a similar level of semantics. To increase the diversity for autoregression, we randomly generate a vector $\mathbf{R}_{ij} = \mathcal{U}(0, 1)$ (where \mathcal{U} denotes the uniform distribution), and add it to the distance \mathbf{D} and obtain the centerness \mathbf{C} :

$$\mathbf{C}_{ij} = \mathbf{D}_{ij} + \lambda \mathbf{R}_{ij} \quad (5)$$

where $\lambda = 0.01$ is set to avoid the randomness change the order for patches with different distances. The centerness \mathbf{C} can be reshaped to one dimension $\mathbf{C} \in \mathbb{R}^N$. Finally, the autoregression permutation can be calculated:

$$\pi = \text{arg sort}(\mathbf{C}) \quad (6)$$

Thus, we generated the semantic-aware autoregressive permutation for image patches.

Autoregressive Modeling Based on the generated autoregressive permutation π , the autoregressive modeling procedure can be conducted. Following previous work (Qi et al. 2022), we design a parallel encoder-decoder architecture to perform autoregressive modeling. During pre-training, the encoder focuses on learning contextual information, and the decoder focuses on predicting the given target of the original image from the latent representation. During fine-tuning, only the encoder will be reserved and fine-tuned for downstream tasks.

Encoder. The encoder learns contextual information with masked self-attention. It has the same structure as the Vision Transformer (Dosovitskiy et al. 2020), which consists of L layers of self-attention blocks. We apply a mask to the self-attention blocks to make the current token see only the preceding tokens in the permutation π . Specifically, the mask is generated as follows:

$$\mathbf{M}_{ij} = \begin{cases} 0, & \mathbf{C}_i < \mathbf{C}_j \\ 1, & \mathbf{C}_i \geq \mathbf{C}_j \end{cases} \quad (7)$$

where $\mathbf{M} \in \mathbb{R}^{N \times N}$, i, j are the coordinate of attention matrix. $\mathbf{M}_{ij} = 1$ represents the i -th token have access to the i -th token, and vice versa. We define the output of l -th encoder layer as $\mathbf{h}_i^{(l)}$, where i is the token index. And the initialized sequence \mathbf{x} is the input of the first encoder layer, i.e. $\mathbf{h}_i^{(0)} = \mathbf{x}_i$. Omitting the layer norm, MLP, and residual connection for simplification, the forward process of the encoder can be described as follows:

$$\mathbf{h}_{\pi_i}^{(l)} = \text{Attention}(\mathbf{Q} = \mathbf{h}_{\pi_i}^{(l-1)}, \mathbf{KV} = \mathbf{h}_{\pi_{<i}}^{(l-1)}; \theta_e^{(m)}) \quad (8)$$

where $1 \leq l \leq L$, $\theta_e^{(l)}$ is the parameters of the l -th encoder layer. Note that the masking strategy is implemented by adding a minus infinity value to the self-attention score where $\mathbf{M}_{ij} = 0$. Thus, during encoding, the current token can only see the preceding tokens.

Decoder. The decoder consists of L' layers of cross-attention blocks and an MLP head. The blocks decode the input signals from the latent representation and the MLP head projects the signal to the dimension of the given target of the original image. A mask is applied to the cross-attention blocks:

$$\mathbf{M}'_{ij} = \begin{cases} 0, & \mathbf{C}_i \leq \mathbf{C}_j \\ 1, & \mathbf{C}_i > \mathbf{C}_j \end{cases} \quad (9)$$

Compared with the encoder mask in Eq. (7), the current token can see the preceding tokens and itself. We can define the output of the l -th decoder layer as $\mathbf{g}_i^{(l)}$. And the position embeddings \mathbf{p} are used as the input of the first decoder layer, i.e., $\mathbf{g}_i^{(0)} = \mathbf{p}_i$. The forward process of the decoder blocks can be formulated as follows:

$$\mathbf{g}_{\pi_i}^{(l)} = \text{Attention}(\mathbf{Q} = \mathbf{g}_{\pi_i}^{(l-1)}, \mathbf{KV} = \mathbf{h}_{\pi_{<i}}^{(l-1)}; \theta_d^{(m)}) \quad (10)$$

where $1 \leq l \leq L'$, $\theta_d^{(l)}$ is the parameters of the l -th decoder layer. Finally, the MLP head projects the output of the decoder blocks to the target space:

$$\mathbf{g}_{\pi_i} = \text{MLP}(\mathbf{g}_{\pi_i}^{(L')}; \theta_h) \quad (11)$$

where θ_h is the parameters of the MLP head.

Prediction Targets. The autoregression model is trained by minimizing the mean squared error between the output prediction \mathbf{g}_{π_i} and given targets $\hat{\mathbf{g}}_{\pi_i}$:

$$\mathcal{L} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \sum_{i=1}^N \|\mathbf{g}_{\pi_i} - \hat{\mathbf{g}}_{\pi_i}\|^2 \quad (12)$$

In previous autoregression image modeling (Qi et al. 2022), the raw pixels of patches are employed as the prediction targets $\hat{\mathbf{g}}$. In this study, considering that the raw pixels of patches are low-level signals and are not ideal prediction targets for learning high-level semantic representation, we also explore utilizing the patch features as the prediction targets. Specifically, we also use the features extracted from the pre-trained ViT-B model of DINO (Caron et al. 2021) and CLIP (Radford et al. 2021) as the prediction targets.

Experiments

Settings

Pre-training. All self-supervised pre-training is performed on the ImageNet (Russakovsky et al. 2015) training set with a resolution of 224×224 . In default settings, we take ViT-B/16 (Dosovitskiy et al. 2020) as the default backbone and pre-train models with a 2048 batch size for 200 epochs. The decoder is a stack of cross-attention Transformer (Dosovitskiy et al. 2020) blocks and has 12 blocks. Details can be found in *Supplementary Material*.

ImageNet classification. After pre-training, we perform end-to-end fine-tuning, linear probing, and k -NN classification on ImageNet (Russakovsky et al. 2015) to evaluate our SemAIM. For fine-tuning, 100 epochs with a 1024 batch size are performed following common practices (He et al. 2022; Bao, Dong, and Wei 2022) by default. For linear probing, 90 epochs with a 4096 batch size are performed following common practices (He et al. 2022; Qi et al. 2022). The implementation of k -NN classification is based on DINO (Caron et al. 2021). We report top-1 accuracy of a single 224×224 resolution on the ImageNet (Russakovsky et al. 2015) validation set.

COCO object detection and instance segmentation. Following previous methods (He et al. 2022; Bao, Dong, and Wei 2022), the Mask R-CNN (He et al. 2017) with FPN (Lin

order	ImageNet fine-tune	COCO detection AP ^b	COCO segmentation AP ^m	ADE20k mIoU
raster	82.9	42.0	37.9	42.2
stochastic	83.0	42.1	38.0	42.4
similarity	82.4	38.9	35.4	41.6
semantic-aware	83.5	43.3	39.1	43.9

Table 1: Ablation study of autoregression order. The RGB value is adopted as the prediction target, and the decoder depth is 12. Note that “similarity” denotes the similarity is directly used as autoregression order.

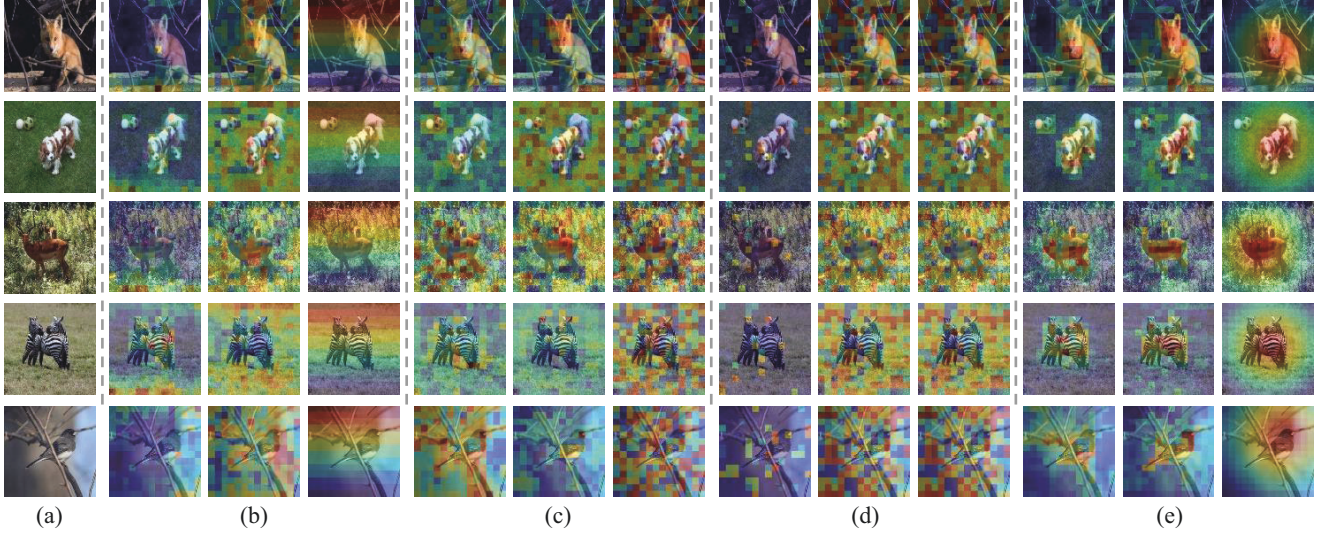


Figure 2: Visualization of different autoregression orders. (a) input images, (b) raster order used in iGPT (Chen et al. 2020a), (c) stochastic order used in SAIM (Qi et al. 2022), (d) similarity order (the similarity map S is also directly used as the autoregression order), and (e) semantic-aware order used in SemAIM. In (b)(c)(d)(e), the first column shows the self-attention maps from the last block, the second column shows similarity maps S from the last block, and the last column shows the corresponding autoregression orders (more warm-colored patches are predicted first).

et al. 2017) is adopted as the detector. We conduct end-to-end fine-tuning on COCO (Lin et al. 2014) with 1024×1024 resolution. We train 12.5 epochs with a 16 batch size for ablations and 100 epochs with a 64 batch size for fair comparison with other methods. AP^b and AP^m are reported for object detection and instance segmentation, respectively. Our implementation is based on detectron2 (Wu et al. 2019) and ViTDet (Li et al. 2022b).

ADE20k semantic segmentation. Following previous works (He et al. 2022; Bao, Dong, and Wei 2022), UperNet (Xiao et al. 2018) is adopted as the decoder. we perform end-to-end fine-tuning on ADE20k (Zhou et al. 2017) for 160k iterations with 512×512 resolution and a 16 batch size. mIoU (Everingham et al. 2009) is used for evaluation. Our implementation is based on mmsegmentation (Contributors 2020).

Ablation Study

In this part, we conduct ablation studies to analyze the influence of each part of SemAIM. Specifically, we analyze the influence of autoregression order, prediction targets, and decoder depth in Tab. 1, Tab. 2, and Tab. 3, respectively. We

take the ViT-B/16 (Dosovitskiy et al. 2020) pre-trained with 200 epochs on ImageNet (Russakovsky et al. 2015) as the backbone and report the results on ImageNet (Russakovsky et al. 2015) classification, COCO (Lin et al. 2014) object detection and instance segmentation, and ADE20k (Zhou et al. 2017) semantic segmentation. In these experiments, 100 epochs of fine-tuning on ImageNet (Russakovsky et al. 2015), 12.5 epochs of fine-tuning on COCO (Lin et al. 2014), and 160k iterations of fine-tuning on ADE20k (Zhou et al. 2017) are performed. The default settings of our SemAIM are shown in bold in the tables.

Influence of autoregression order. In this experiment, we analyze the influence of autoregression order for representation learning. The raster and stochastic order adopted in iGPT (Chen et al. 2020a) and SAIM (Qi et al. 2022) are compared with the semantic-aware order adopted in SemAIM. In addition, the similarity map S is also directly used as the autoregression order (patch with larger similarity is predicted first), which is denoted as “similarity”. As shown in Tab. 1, the semantic-aware order proposed in this study significantly outperforms the raster and stochastic order, which verifies that semantic-aware prediction is more suitable

targets	ImageNet fine-tune	COCO detection AP ^b	COCO segmentation AP ^m	ADE20k mIoU
RGB	83.5	43.3	39.1	43.9
DINO	84.4	47.8	42.3	48.3
CLIP	84.6	47.6	42.1	51.1

Table 2: Ablation study of prediction targets. The semantic-aware order is adopted, and the decoder depth is 12.

d	ImageNet fine-tune	COCO detection AP ^b	COCO segmentation AP ^m	ADE20k mIoU
1	82.1	42.0	38.2	42.3
6	83.3	43.1	39.0	43.3
12	83.5	43.3	39.1	43.9

Table 3: Ablation study of decoder depth. The semantic-aware order is adopted, and the RGB value is adopted as the prediction target.

for autoregression image modeling. Further, the “similarity” order achieves the worst performance, which indicates that it is noisy and unstable to utilize the similarity map for autoregression order directly.

In addition, we visualize the self-attention maps, and the similarity maps S , and the corresponding autoregression orders of each method in Fig. 2. The self-attention maps and the similarity maps of the semantic-aware order used in SemAIM locate on semantic regions more accurately than other methods. This indicates that SemAIM can learn more semantic representations.

Influence of prediction targets. In this experiment, we analyze the influence of prediction targets. Apart from the RGB value, we also utilize the feature of the pre-trained ViT-B model from DINO (Caron et al. 2021) and CLIP (Radford et al. 2021) as the prediction targets. As shown in Tab. 2, using features as prediction targets perform better than RGB value, which indicates that utilizing high-level features as prediction targets is more beneficial for learning high-level semantic representation.

Influence of decoder depth. In this experiment, we analyze the influence of the decoder depth. Tab. 3 varies the decoder depth L' (number of Transformer blocks). When the decoder depth is less than the encoder depth, i.e., $L' < L$, we uniformly choose L' layer of the encoder to connect with the decoder during the parallel autoregression modeling procedure. As can be seen, a sufficiently deep decoder is essential for the performance of SemAIM. This result demonstrates that autoregression image modeling needs a deeper decoder than masked image modeling (He et al. 2022).

Comparisons with Other Methods

The proposed SemAIM is compared with a wide range of self-supervised counterparts on ImageNet (Russakovsky et al. 2015) classification, COCO (Lin et al. 2014) object detection and instance segmentation, and ADE20k (Zhou et al. 2017) semantic segmentation.

We compare our SemAIM with a wide range of self-supervised methods, including contrastive learning meth-

ods (Caron et al. 2021; Chen, Xie, and He 2021), masked image modeling methods (Bao, Dong, and Wei 2022; He et al. 2022; Xie et al. 2022; Li et al. 2022a; Wang et al. 2023b,a), contrastive learning and masked image modeling combinations (Zhou et al. 2022; Dong et al. 2022), and autoregression image modeling (Chen et al. 2020a; Hua et al. 2022; Qi et al. 2022). All methods are pre-trained with the same resolution, i.e., 224×224 on ImageNet-1K (Russakovsky et al. 2015).

ImageNet classification. We compare the proposed SemAIM with state-of-the-art alternatives on the ImageNet-1K (Russakovsky et al. 2015). The results are shown in Tab. 4. Notably, with only 400 epochs pre-training, SemAIM achieves 83.8% using ViT-B/16 as the backbone, surpassing MAE (He et al. 2022) pre-trained for 1600 epochs by +0.2%. This empirical evidence demonstrates that enhancing the semantic awareness of ViTs during autoregression modeling brings better visual representations. With 800 epochs of pre-training, SemAIM outperforms its baseline SAIM (Qi et al. 2022) by 0.2% using ViT-B/16 as the backbone. It achieves competitive results compared with the state-of-the-art. Specifically, it achieves 84.1% and 85.8% using ViT-B/16 and ViT-L/16, respectively. Furthermore, SemAIM reaches state-of-the-art results, 85.3% and 86.5% using ViT-B/16 and ViT-L/16, respectively, by using CLIP (Radford et al. 2021) feature as predict targets.

COCO object detection and segmentation. Following the configuration of ViTDet (Li et al. 2022b), the pre-trained models are fine-tuned on COCO (Lin et al. 2014) with 100 epochs, using the Mask R-CNN (He et al. 2017) detector. We take ViT-B/16 (Dosovitskiy et al. 2020) as the backbone. AP^b and AP^m are adopted as the metric for object detection and instance segmentation, respectively. As shown in Tab. 5, with only 400 epochs of pre-training, SemAIM achieves 50.7% AP^b and 45.0% AP^m, outperforming the baseline SAIM (Qi et al. 2022) pre-trained for 800 epochs by +1.3% and +1.0%, respectively. With 800 epochs of pre-training, SemAIM achieves 51.3% AP^b and 45.4% AP^m, surpassing SAIM (Qi et al. 2022) by +1.9% and +1.4%, respectively. The

method	ep.	ViT-B	ViT-L
<i>Contrastive Learning</i>			
DINO (Caron et al. 2021)	1600	82.8	-
MoCo v3 (Chen, Xie, and He 2021)	600	83.2	84.1
<i>Masked Image Modeling</i>			
BEiT (Bao, Dong, and Wei 2022)	800	83.2	85.2
MAE (He et al. 2022)	1600	83.6	85.9
SimMIM (Xie et al. 2022)	800	83.8	-
SemMAE (Li et al. 2022a)	800	83.4	-
LocalMIM (Wang et al. 2023b)	1600	84.0	-
HPM (Wang et al. 2023a)	800	84.2	85.8
<i>Masked Image Modeling + Contrastive Learning</i>			
iBOT (Zhou et al. 2022)	1600	84.0	-
BootMAE (Dong et al. 2022)	800	84.2	85.9
<i>Autoregression Image Modeling</i>			
iGPT (Chen et al. 2020a)	-	72.6 [‡]	-
ViT-iGPT (Chen et al. 2020a)	-	82.7 [‡]	-
RandSAC (Hua et al. 2022)	1600	83.9	-
SAIM (Qi et al. 2022)	800	83.9	-
SemAIM	400	83.8	85.5
SemAIM	800	84.1	85.8
SemAIM[†]	800	85.3	86.5

Table 4: Comparison with previous methods on ImageNet-1K classification. All methods are evaluated by fine-tuning. The resolution of images is 224×224 for both pre-training and fine-tuning. [†] means using CLIP (Radford et al. 2021) feature as predict targets. [‡] means the result is borrowed from (Qi et al. 2022).

improvement is more significant than those on the ImageNet classification shown in Tab. 4, indicating that semantic-aware autoregression modeling can learn better spatial reasoning abilities.

ADE20k semantic segmentation. The pre-trained models are fine-tuned on ADE20k (Zhou et al. 2017) for semantic segmentation with 160k iterations using UperNet (Xiao et al. 2018) We take ViT-B/16 (Dosovitskiy et al. 2020) as the backbone and search for the optimal learning rate for each model. mIoU is used as the metric. As shown in Tab. 5, with 800 epochs of pre-training, SemAIM achieves 48.0% mIoU, surpassing the baseline SAIM (Qi et al. 2022) by +0.2%.

Conclusion

In this study, we present a semantic-aware autoregressive image modeling (SemAIM) method for visual representation learning. The key insight of SemAIM is to model images from the most semantic patches to the less semantic patches autoregressively. In SemAIM, we first calculate a semantic-aware permutation of patches according to their feature similarities and then perform the autoregression procedure based on the permutation. In addition, considering that the raw pixels of patches are low-level signals and are not ideal prediction targets for learning high-level semantic representation, we also explore utilizing the patch features as

method	ep.	COCO AP ^b	AP ^m	ADE20k mIoU
supervised				
MoCo v3 (Chen, Xie, and He 2021)	600	47.9 [†]	42.7 [†]	47.3 [†]
BEiT (Bao, Dong, and Wei 2022)	800	49.8 [†]	44.4 [†]	47.1
MAE (He et al. 2022)	1600	50.3	44.9	48.1
CAE (Chen et al. 2022)	1600	50.1	44.0	48.8
PeCo (Dong et al. 2023)	1600	49.1	43.8	48.5
SIM (Tao et al. 2023)	800	49.1	43.8	-
HPM (Wang et al. 2023a)	800	50.1	44.6	48.5
SAIM (Qi et al. 2022)	800	49.4	44.0	47.8
SemAIM	400	50.7	45.0	47.7
SemAIM	800	51.3	45.4	48.0

Table 5: Comparison with other methods on downstream tasks. All methods take the ViT-B/16 (Dosovitskiy et al. 2020) as the backbone. For COCO (Lin et al. 2014) object detection and instance segmentation, we utilize Mask R-CNN (He et al. 2017) and perform 100 epoch fine-tuning. For ADE20k (Zhou et al. 2017) semantic segmentation, we use UperNet (Xiao et al. 2018) and perform 160k iterations of fine-tuning. [†] means the result is borrowed from (He et al. 2022).

the prediction targets. We conducted extensive experiments on a broad range of downstream tasks, including image classification, object detection, and instance/semantic segmentation, to evaluate the performance of SemAIM. The results demonstrate SemAIM achieves state-of-the-art performance compared with other self-supervised methods. This study demonstrates that it is crucial for autoregressive image modeling to perform a suitable, semantic-aware autoregressive permutation. SemAIM shows good performance for vision pre-training, and it is a unified pre-training task compared with language modeling in NLP. One potential limitation of SemAIM is that we only consider the case of one center patch in the permutation generation process, which may influence the performance of images with multiple objects. This limitation can be overcome by calculating multiple center patches in the permutation generation process.

References

- Bao, H.; Dong, L.; and Wei, F. 2022. Beit: Bert pre-training of image transformers. In *International Conference on Learning Representations (ICLR)*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems (NeurIPS)*, 33: 9912–9924.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of*

- the *IEEE/CVF International Conference on Computer Vision (ICCV)*, 9650–9660.
- Chen, M.; Radford, A.; Child, R.; Wu, J.; Jun, H.; Luan, D.; and Sutskever, I. 2020a. Generative pretraining from pixels. In *International conference on machine learning*, 1691–1703. PMLR.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020b. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, T.; Saxena, S.; Li, L.; Fleet, D. J.; and Hinton, G. 2021. Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852*.
- Chen, X.; Ding, M.; Wang, X.; Xin, Y.; Mo, S.; Wang, Y.; Han, S.; Luo, P.; Zeng, G.; and Wang, J. 2022. Context autoencoder for self-supervised representation learning. *arXiv preprint arXiv:2202.03026*.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15750–15758.
- Chen, X.; Xie, S.; and He, K. 2021. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9640–9649.
- Contributors, M. 2020. MMSegmentation: OpenMMLab Semantic Segmentation Toolbox and Benchmark. <https://github.com/open-mmlab/mms Segmentation>.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dong, L.; Yang, N.; Wang, W.; Wei, F.; Liu, X.; Wang, Y.; Gao, J.; Zhou, M.; and Hon, H.-W. 2019. Unified language model pre-training for natural language understanding and generation. *Advances in neural information processing systems*, 32.
- Dong, X.; Bao, J.; Zhang, T.; Chen, D.; Zhang, W.; Yuan, L.; Chen, D.; Wen, F.; and Yu, N. 2022. Bootstrapped Masked Autoencoders for Vision BERT Pretraining. In *European Conference on Computer Vision (ECCV)*, 247–264. Springer.
- Dong, X.; Bao, J.; Zhang, T.; Chen, D.; Zhang, W.; Yuan, L.; Chen, D.; Wen, F.; and Yu, N. 2023. Peco: Perceptual codebook for bert pre-training of vision transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Eriksen, C. W.; and Hoffman, J. E. 1972. Temporal and spatial characteristics of selective encoding from visual displays. *Perception & psychophysics*, 12: 201–204.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2009. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision (IJCV)*, 88: 303–308.
- Guo, Y.; Xu, M.; Li, J.; Ni, B.; Zhu, X.; Sun, Z.; and Xu, Y. 2022. HCSC: Hierarchical Contrastive Selective Coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9706–9715.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16000–16009.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9729–9738.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2961–2969.
- Hou, Z.; Sun, F.; Chen, Y.-K.; Xie, Y.; and Kung, S.-Y. 2023. Milan: Masked image pretraining on language assisted representation.
- Hua, T.; Tian, Y.; Ren, S.; Zhao, H.; and Sigal, L. 2022. Self-supervision through random segments with autoregressive coding (randsac). *International Conference on Learning Representations (ICLR)*.
- Kakogeorgiou, I.; Gidaris, S.; Psomas, B.; Avrithis, Y.; Bursuc, A.; Karantzalos, K.; and Komodakis, N. 2022. What to hide from your students: Attention-guided masked image modeling. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXX*, 300–318. Springer.
- Koch, K.; McLean, J.; Segev, R.; Freed, M. A.; Berry, M. J.; Balasubramanian, V.; and Sterling, P. 2006. How much the eye tells the brain. *Current biology*, 16(14): 1428–1434.
- Komodakis, N.; and Gidaris, S. 2018. Unsupervised representation learning by predicting image rotations. In *International conference on learning representations (ICLR)*.
- Li, G.; Zheng, H.; Liu, D.; Wang, C.; Su, B.; and Zheng, C. 2022a. SemMAE: Semantic-Guided Masking for Learning Masked Autoencoders. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Li, Y.; Mao, H.; Girshick, R.; and He, K. 2022b. Exploring plain vision transformer backbones for object detection. In *European Conference on Computer Vision (ECCV)*, 280–296. Springer.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2117–2125.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 740–755. Springer.
- Noroozi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VI*, 69–84. Springer.

- Noroozi, M.; Pirsiavash, H.; and Favaro, P. 2017. Representation learning by learning to count. In *Proceedings of the IEEE international conference on computer vision*, 5898–5906.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Pathak, D.; Krahenbuhl, P.; Donahue, J.; Darrell, T.; and Efros, A. A. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2536–2544.
- Qi, Y.; Yang, F.; Zhu, Y.; Liu, Y.; Wu, L.; Zhao, R.; and Li, W. 2022. Exploring Stochastic Autoregressive Image Modeling for Visual Representation. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, 8748–8763. PMLR.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, 8821–8831. PMLR.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115: 211–252.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2020. Mpnnet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33: 16857–16867.
- Song, K.; Xie, J.; Zhang, S.; and Luo, Z. 2023a. Multi-Mode Online Knowledge Distillation for Self-Supervised Visual Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11848–11857.
- Song, K.; Zhang, S.; Luo, Z.; Wang, T.; and Xie, J. 2023b. Semantics-Consistent Feature Search for Self-Supervised Visual Representation Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 16099–16108.
- Tao, C.; Zhu, X.; Su, W.; Huang, G.; Li, B.; Zhou, J.; Qiao, Y.; Wang, X.; and Dai, J. 2023. Siamese image modeling for self-supervised vision representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2132–2141.
- Van den Oord, A.; Kalchbrenner, N.; Espeholt, L.; Vinyals, O.; Graves, A.; et al. 2016. Conditional image generation with pixelcnn decoders. *Advances in neural information processing systems*, 29.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.
- Wang, H.; Song, K.; Fan, J.; Wang, Y.; Xie, J.; and Zhang, Z. 2023a. Hard Patches Mining for Masked Image Modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wang, H.; Tang, Y.; Wang, Y.; Guo, J.; Deng, Z.-H.; and Han, K. 2023b. Masked Image Modeling with Local Multi-Scale Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wei, C.; Fan, H.; Xie, S.; Wu, C.-Y.; Yuille, A.; and Feichtenhofer, C. 2022. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14668–14678.
- Wu, Y.; Kirillov, A.; Massa, F.; Lo, W.-Y.; and Girshick, R. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3733–3742.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified perceptual parsing for scene understanding. In *European Conference on Computer Vision (ECCV)*, 418–434.
- Xie, Z.; Zhang, Z.; Cao, Y.; Lin, Y.; Bao, J.; Yao, Z.; Dai, Q.; and Hu, H. 2022. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9653–9663.
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R. R.; and Le, Q. V. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, 649–666. Springer.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 633–641.
- Zhou, J.; Wei, C.; Wang, H.; Shen, W.; Xie, C.; Yuille, A.; and Kong, T. 2022. Image BERT Pre-training with Online Tokenizer. In *International Conference on Learning Representations (ICLR)*.