

Evolved Hierarchical Masking for Self-Supervised Learning

Zhanzhou Feng, Shiliang Zhang *Senior Member, IEEE*

Abstract—Existing Masked Image Modeling methods apply fixed mask patterns to guide the self-supervised training. As those mask patterns resort to different criteria to depict image contents, sticking to a fixed pattern leads to a limited vision cues modeling capability. This paper introduces an evolved hierarchical masking method to pursue general visual cues modeling in self-supervised learning. The proposed method leverages the vision model being trained to parse the input visual cues into a hierarchy structure, which is hence adopted to generate masks accordingly. The accuracy of hierarchy is on par with the capability of the model being trained, leading to evolved mask patterns at different training stages. Initially, generated masks focus on low-level visual cues to grasp basic textures, then gradually evolve to depict higher-level cues to reinforce the learning of more complicated object semantics and contexts. Our method does not require extra pre-trained models or annotations and ensures training efficiency by evolving the training difficulty. We conduct extensive experiments on seven downstream tasks including partial-duplicate image retrieval relying on low-level details, as well as image classification and semantic segmentation that require semantic parsing capability. Experimental results demonstrate that it substantially boosts performance across these tasks. For instance, it surpasses the recent MAE by 1.1% in imageNet-1K classification and 1.4% in ADE20K segmentation with the same training epochs. We also align the proposed method with the current research focus on LLMs. The proposed approach bridges the gap with large-scale pre-training on semantic demanding tasks and enhances intricate detail perception in tasks requiring low-level feature recognition.

Index Terms—Self-supervised learning, Masked image modeling, Efficient learning, Model pretraining.

I. INTRODUCTION

RECENT years have witnessed a boom in continuously growing representation learning capability and data demands of deep neural networks. To tackle the increasing demand for labeled data, Self-Supervised Learning (SSL) has emerged as a promising paradigm in pre-training deep models. By constructing pretext training tasks, the pretrained model can acquire transferable representations from unlabeled large-scale data. In the field of natural language processing, Masked Language Modeling (MLM) has emerged as a predominant pre-training task, constituting an essential component in the training process of renowned models such as BERT [1] and GPT [2]. MLM masks several words in the input sentences and supervises the network to recover masked words according to the semantics provided by the remaining words.

Zhanzhou Feng and Shiliang Zhang are with National Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University, Beijing 100871, China (e-mail: fengzz@stu.pku.edu.cn; slzhang.jdl@pku.edu.cn).

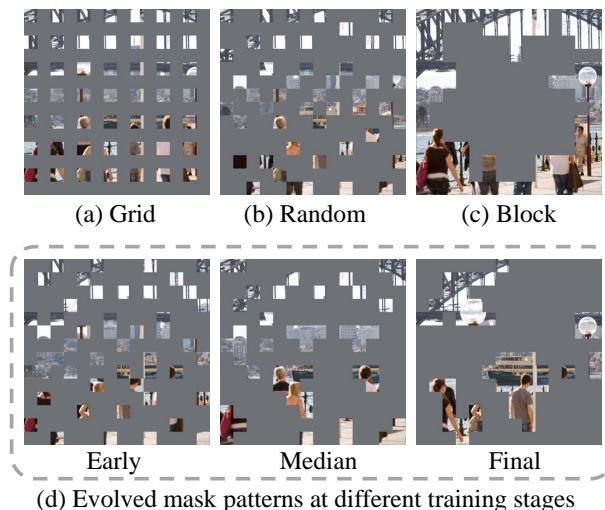


Fig. 1. (a), (b), and (c) are three basic mask patterns adopted in existing MIM methods. (d) illustrates the proposed evolved hierarchical masking, where the generated mask patterns evolve with the capability of the vision model being trained.

Inspired by the success of MLM, a vision counterpart, Masked Image Modeling (MIM) is proposed to pre-train vision models on unlabeled images. MIM follows a similar idea from language to mask a portion of input image patches, then the pretext task is to recover masked contents from visible patches. As images are not structured like sentences, different MIM works have to resort to different criteria to generate mask patterns. We categorize mask patterns in existing works into three based on their followed masking criteria. Some works like MAE [3] and SimMIM [4] randomly mask image patches, assuming input information density is uniformly distributed. Another line of works, like MST [5], aims to preserve crucial cues in the image to enhance the learning of local context. A third line, including AttnMask [6], I-JEPA [7], and Sem-MAE [8], suggests completely masking cues such as object regions in images to pose a more challenging learning task.

To study the impact of mask patterns on self-supervised pre-training, we test the performance of three commonly used masking methods, i.e., random masks, grid masks and block masks, as shown in Fig. 1(a)-(c). Fig. 2(a) explores their effects on two vision tasks. It can be observed that, random pattern and block pattern perform best in image classification and semantic segmentation, respectively. It is also clear that, more training epochs do not boost the performance of grid pattern and random pattern in segmentation. Fig. 2(b) further

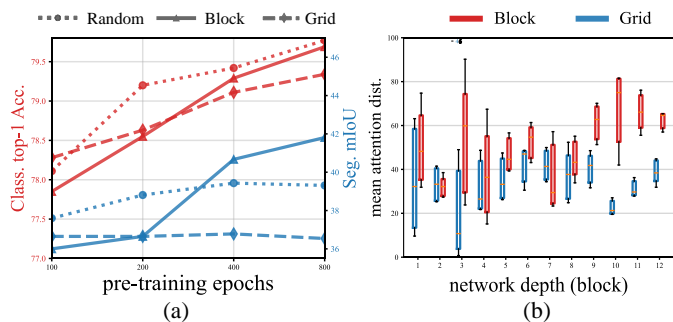



Fig. 2. Illustration of effects of different mask patterns to downstream tasks in (a), and learned parameters in (b), respectively. In (a), random pattern and block pattern perform best in image classification and semantic segmentation, respectively. (b) shows the mean attention distance across images at different layers of the pre-trained model. The unit of the y-axis is the pixel. Those results indicate different mask patterns are suited to different tasks.

visualizes the average attention length of neurons at each layer of the pre-trained model. We compute attention distance by averaging the distance between the query tokens and all other tokens, weighted by the attention weight following [9]. It indicates that, neurons trained by grid mask mostly focus on nearby regions with shorter attention distances. As longer attention distance benefits the learning of contextual cues, block pattern is more preferred by dense prediction tasks like semantic segmentation.


The above experiments indicate that, the choice of patterns for generating masks largely determines visual dependencies that the network could learn in pre-training. For instance, masking the entire object regions enforces the network to learn correlated semantics and contexts of visible cues, which are helpful for image classification and object segmentation. On the other hand, grid mask patterns supervises the network neuron to attend to nearby regions, which are important for recovering low-level details and enhancing efficiency in learning tasks. Therefore, different mask patterns are suitable for different downstream tasks. This poses a challenge for self-supervised learning, i.e., the pre-training procedure has no clue which task it will be applied to. In other words, visual knowledge learned with fixed masking patterns is limited and hard to generalize to various downstream tasks.

Mask pattern adopted by MLM is hence expected to enforce the learning with various visual cues. We propose to parse image contents into a hierarchy representing different levels of visual cues, and progressively mask visual cues according to this hierarchy. An adaptive part partition module is adopted to leverage the vision model being trained to construct the hierarchy, and generate mask patterns. Establishment of the hierarchy starts by treating each image patch as a leaf node in a tree. We take the attention map produced by the current model as the similarity matrix among image patches. A pair of most similar nodes are hence merged to generate their parent node. This procedure continues until a single root encompasses all leaf nodes, which leads to a structured hierarchy representing visual cues at various scales. It's worth noting that the hierarchy is dynamically constructed each time an image is input into the system; thus, it evolves alongside the model training.

High-level Cues



Middle-level Cues



Low-level Cues



Fig. 3. Visualization of [CLS] attention heatmap from classification [10] in (a), semantic segmentation result [11] in (b), and landmark retrieval results on Oxford Building [12], [13] in (c). These visualizations show the proposed method exhibits superior visual cue acquisition capacity at different semantic levels. Best viewing with zoom-in.

As the model progressively learns new visual information, the generated hierarchical structure also improves.

We then perform masking by removing image patches on a certain level of the established hierarchy. Initially, masked patches are sampled from low-level hierarchies to depict visual details. This enforces the network to learn low-level visual cues, which are prerequisites for developing higher-level representations. As the training progresses, the masked location gradually advances to higher hierarchies to depict more complicated regions as illustrated in Fig. 1(d). In other words, this transition from low to high hierarchy generates simple patterns at the initial training stage, which hence evolves to eliminate accurate object parts to reinforce the learning of object semantics and contexts. This procedure progressively masks visual cues and reveals the underlying visual dependencies, e.g., from local textures to object parts and general semantics, to gradually promote the representation learning capabilities.

We test the proposed method on various downstream tasks requiring different visual cues modeling capabilities. As illustrated in Fig. 3 (a), the attention map of our method focuses on both the foreground player and the background referee, indicating a more comprehensive understanding of the boxing match scene. Fig. 3(b) shows that our method performs better than previous works in image segmentation, showing its boosted performance in inferring region and pixel-level semantics. Fig. 3(c) presents the results of image matching, which heavily relies on the discriminative power of local features extracted at corner points. It is clear that, our method gets the most reliable feature matching result among competitors.

Extensive experimental results also demonstrate that the proposed method produces a stronger initialization model for various downstream tasks. We test the effectiveness of the

proposed method on three popular MIM architectures, *i.e.*, MAE [3], BEiT [14] and SimMIM [4]. Our method brings significant performance enhancement for those three architectures, especially on the semantic segmentation task, *e.g.*, boosts the mIoU by +1.42%, +1.39% and +1.32%, respectively. When compared with recent self-supervised learning methods, our method achieves comparable performance with fewer pre-training epochs, and superior performance with similar training epochs. When combined the proposed method with Large Language Models, it bridges the performance gap between self-supervised techniques and large-scale pre-training in semantic demanding VQA_{v2} [15] task by +2.2%. Additionally, it amplifies the strength of self-supervised methods in recognizing intricate details.

To the best of our knowledge, this is an original effort on evolved hierarchy masking for self-supervised learning. It builds a semantic hierarchy to generate masks, leading to an evolved masking criteria at different training stages. Our method does not require extra pre-trained models or annotations. It effectively ensures training efficiency, and enhances the generalization ability of the trained model. As shown in various downstream tasks, our method shows better potentials to boost the performance of pre-trained vision models.

II. RELATED WORKS

This work is closely related to self-supervised learning and masked image modeling. This section briefly reviews related works in recent years.

Self-supervised Learning. Generally, SSL involves creating an annotation-free pretext task that helps pre-trained models learn visual representations by understanding connections in visual inputs. Existing SSL methods can be divided into two categories based on how annotation-free tasks are designed.

The first category, contrastive learning, has dominated the learning algorithms in those works for the past few years. Under an instance discrimination setting, each image is considered as an individual class [16]. It works by pulling positive samples together and pushing negative samples apart. A common technique involves optimizing an encoder to generate comparable embeddings for multiple perspectives of the same image. These perspectives are usually created using a range of carefully designed data augmentations, including random cropping [17], horizontal flipping [18], and color adjustments [19], among others [20]. However, this technique requires prior knowledge of task-specific invariances, which is often unknown during the pretraining phase. Related works includes SimCLR [17], MoCo [21], Swav [22], and BYOL [23], *etc.* DINO [24] formulates self-supervised contrastive as self-distillation. The student model is trained to predict the output of a teacher network built with a momentum encoder.

Another category of SSL research predicts the original image based on the partially observed data. For example, RotNet [25] predicts the 2D rotation applied to the input image. The CFN [26] randomly shuffles the image patches and takes Jigsaw puzzles as the pretext task. Autoencoder is a commonly used generative SSL model, which is trained by minimizing the reconstruction error. It has an encoder

that maps the input data to a latent space and a decoder to reconstruct the image conditioned on the latent representation. Denoising autoencoders (DAE) [27] corrupts input signal with random noise and context encoders regress an entire image region based on its surroundings. A series of methods can be viewed as generalized DAE with different ways of generating corrupted images, including degrading the resolution [28], masking regions [29], or removing certain color channels [30], *etc.* MIM also can be regarded as one of DAE variants.

Masked Image Modeling. Inspired by the success of MLM [1], [2] in NLP, MIM has been proposed to tackle the data-hungry issue of vision transformers [9], [31]. As one of the core designs of MIM, the mask methods largely determine the knowledge that the network acquires in the pre-training phase. According to the masking criteria, this work divides existing MIM methods into three categories and discusses their characteristics and trade-offs.

The first category involves random masking, a widely adopted and straightforward approach. MAE [3] is one of the representative works that utilize an asymmetric autoencoder to recover a randomly masked input. SimMIM [4] randomly masks larger square patches and minimizes the ℓ_1 loss between raw pixel values and predicted results. The advantage of these methods lies in their simplicity and ease of implementation. However, due to the complete randomness of the generated masks, extensive training iterations are required to ensure a comprehensive coverage of visual dependencies conveyed from the generated masks. This inefficiency further increases the burden of computing resources during pre-training.

The second category prioritizes reserving crucial cues within the masking process. For example, MST [5] masks only nonessential patches and preserves key patterns in images. MFM [32] uses low-pass/high-pass filters to perform masking, and most object regions with clear semantics are preserved. These works preserve the most important contents in the image, posing an easier MIM task. They facilitate rapid and stable convergence, enhance training efficiency, and effectively capture low-level visual details. The limitation of this masking pattern lies in excessive shared information between masked and visible patches. This can lead the network to mimic visible patches for content reconstruction, rather than developing high-level semantic understanding, resulting in subpar high-level semantic learning.

The third category proposes to mask clues like object regions completely. BEiT [14] employs a block-wise masking method to mask some image objects as a whole. AttnMask [6] proposes to mask patches belonging to the most attended objects. SemMAE [8] leverages the iBOT [33] for semantic segmentation and produces the mask according to the segmentation result. Besides, ADIOS [34] utilizes a learned adversarial masking subnet to pose a more challenging MIM task. Mahmoud's proposed masking strategy in I-JEPA [7] emphasizes the importance of selecting target blocks of sufficiently large scale and using a context block that provides ample information. This category excels in modeling both high-level semantics, yielding top performance in several downstream tasks. However, as most pre-trained networks start with randomly initialized parameters, there is a deficiency in

modeling low-level visual cues. In the absence of essential low-level details such as textures and edges, the deduced high-level semantics are more difficult to learn, resulting in convergence challenges and even potential training collapses. Compared to SemMAE [8] and ADIOS [34], our method does not introduce extra networks or training costs. It leverages the model being trained to determine the cues that should be masked. Evolved masks make training difficulty on par with the capability of the network being trained, hence ensuring a more effective and fast self-supervised learning.

Multimodal LLMs. Large Language Models (LLMs) [35], [36] have demonstrated remarkable performance in various tasks without task-specific fine-tuning. Recently, researchers have explored combining LLMs with visual encoders to tackle multi-modal and even pure visual tasks, such as MiniGPT-4 [37] and BLIP-2 [38]. Under this framework, a visual tokenizer encodes the image, followed by bridges like MLP or Perceiver Resampler to map to linguistic space for the LLMs to understand. When properly fed visual data, LLMs comprehend the visual realm and respond to instructions. Therefore, the quality of visual cues extracted by the visual encoder directly impacts the LLMs' inference quality. There are three popular categories of visual encoders: fully supervised, text-guided weakly supervised, and self-supervised. The common fully supervised method [31] involves pretraining on a well-annotated dataset like ImageNet. Text-guided weakly supervised models pre-train on large image-text pairs using contrastive learning. A notable model, CLIP [39], excels in various tasks without specific training. Our method falls under self-supervised pre-training [3], [24].

Difference with previous works. This work is most related to existing masked image modeling works. Different from static masking patterns in existing works, the proposed method introduces evolved masks to consistently gain the model with low, mid, and high-level visual cues modeling capabilities. Those evolved masks can be easily generated thanks to the organized hierarchy representing multi-levels of image contents. As shown in experiments, the proposed method outperforms existing works by clear margins across different tasks. Our previous conference version [40] reveals that specific mask patterns, such as grid-like and large block masks, enhance the learning of distinct visual knowledge. We thus propose the evolved masking strategy by combining different mask patterns. However, simply combining mask patterns makes the proposed method inefficient and hard to tune due to considerable hyper-parameters. This updated version thus introduces more principled methodology by formulating the evolved masking as a hierarchical masking procedure that employs a unified mask generation process controlled by a single hyper-parameter, i.e., the masking depth. It enables a smooth transition between masking low-level and higher-level visual cues, leading to substantial performance enhancement. We also align the proposed method with the current research on LLMs and evaluate it across various tasks requiring different levels of visual cues, including Visual Question Answering, Object Counting, Multi-Class Identification, partial-duplicated retrieval, etc. Extensive evaluations further show the promising performance of the evolved masking strategy.

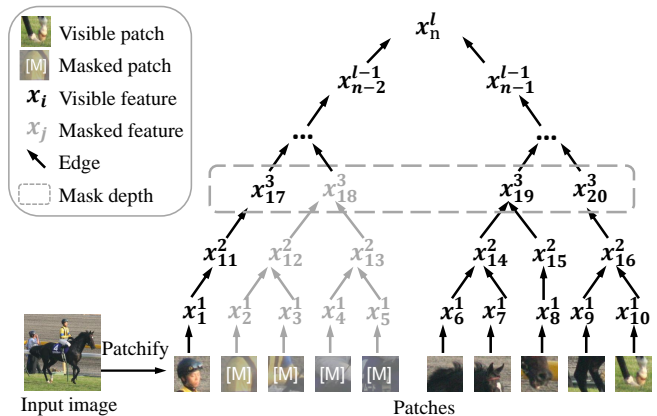


Fig. 4. Illustration of established visual cue hierarchy structure \mathcal{T} . The MIM is performed by reconstructing masked patches according to visible regions. Different mask patterns can be adopted by alternating the mask depth on the hierarchy structure, e.g., masking nodes on the 3^{rd} level or the 1^{st} level.

III. FORMULATION

Given an unlabeled image dataset $\mathcal{D} = \{I\}$, the goal of masking image modeling is to endow pre-trained model with the ability to extract distinguishable features from visual signals. For an input image $I \in \mathbb{R}^{HW \times 3}$, where H, W are the spatial size, we generate a binary mask $\mathcal{M} \in \{0, 1\}^{HW}$ on I , and apply \mathcal{M} to self-supervised learning. Denote by $0 < r < 1$ the masked ratio, controlling number of masked patches, that

$$\sum_{i=0}^{HW} \mathcal{M}_i = HW \times r. \quad (1)$$

Specifically, we adopt an encoder-decoder structure with learnable parameters θ, θ' to recover I from a masked input. For the t -th training epoch, the training objective is denoted as,

$$\arg \min_{\theta, \theta'} \mathbb{E}_{I \sim \mathcal{D}} \mathcal{H}(G_{\theta'}^{(t)}(F_{\theta}^{(t)}(I \odot \mathcal{M})), I \odot (1 - \mathcal{M})), \quad (2)$$

where \odot is the element-wise product, $I \odot \mathcal{M}$ denotes the masked input. $F_{\theta}^{(t)}$ and $G_{\theta'}^{(t)}$ are encoder and decoder in t -th training epoch, respectively. $\mathcal{H}(\cdot, \cdot)$ is the similarity measurement, e.g., l_2 -distance [3] or cross-entropy [14].

Masks \mathcal{M} largely determine encoded cues in the optimized parameters θ and θ' according to Eq (2). In particular, there is a multitude of statistical dependencies among pixels, determined by their spatial location and visual semantics [41]–[43]. The network learns connections between underlying visual cues depending on masked and visible contents. Visual modeling capability of a well-pretrained neural network allows it to transform raw pixels into various levels of abstract concepts for diverse tasks.

To facilitate mask generation, we establish a binary tree structure, $\mathcal{T} := (V, E)$, to represent the relationships between different visual cues, as illustrated in Fig. 4. It consists of a set of latent nodes $V := \{x_i^a\}$, where subscript i represents the sequence index, superscript a denotes layer index. $E := \{e_{i,j}\}$ denotes a set of directed edges, where $e_{i,j}$ connects nodes x_i^a and x_j^b , demonstrating a causal relationship between them. It also indicates the visual cue dependencies the network learned.

To simplify the presentation, we further define the set of parent and children of x_i^a as Pa_i^a and Ch_i^a . We also use De_i^a to denote descendant nodes of node x_i^a . Leaf nodes x_i^1 with superscript $a = 1$ correspond to image patches, where $|\{x_i^1\}| = H \times W$ is the number of image patches. Variables x_i^a in lower layers correspond to intricate details such as texture. Latent variable x_i^a with $a > 1$ carries more complicated semantics. As illustrated in Fig. 4, nodes x_6^1 and x_7^1 depict certain details in the head of horse and node x_8^1 portrays the mouth. The node x_{19}^3 represents the entire horse head.

Constructing the \mathcal{T} requires to model relationship among image patches. MIM pre-training is implemented by inferring masked cues from visible patches. This procedure supervises the network to learn relationships between different regions, hence developing certain visual cue modeling ability as illustrated in previous image representation [44], [45] and generation [46]–[48] works. We hence leverage the visual modeling capability of the pre-trained network to parse input images and construct \mathcal{T} . Specifically, we propose an Adaptive Hierarchy Establishment (AHE) module to leverage features extracted by the encoder $F_\theta^{(t)}$. This can be conceptually denoted as,

$$\mathcal{T}^{(t)} = \text{AHE}(F_\theta^{(t)}, I). \quad (3)$$

Note that, it is easier to model low-level details than learning complicated semantics. We hence progressively construct the \mathcal{T} from layer 1 to layer l at different training epoches to fully leverage the model being pre-trained.

The constructed tree $\mathcal{T}^{(t)}$ is adopted to supervise the MIM at subsequent training epoches. As \mathcal{T} represents visual cues at different levels, different mask patterns can be adopted by alternating masking depths. As illustrated in Fig. 4, we can set masking depth to $a = 3$, and generate masks by randomly selecting nodes x_{18}^3 . Descendant leaf nodes of x_{18}^3 will be masked. The network is hence optimized towards inferring the masked latent variable x_{18}^3 according to other visible patches. Setting a lower mask depth to $a = 1$ would pose easier MIM tasks of modeling low-level details. This procedure enables the MIM task to adaptively generate mask patterns at different training stages.

We hence leverage $\mathcal{T}^{(t)}$ to evolve the mask patterns at different training stages to enhance the generalization capability to different visual cues. Simple masks are generated to learn low-level visual cues at the initial training stage. More complicated masks are hence generated to learn object semantics at a later training stage. We use the training epoch t to represent different training stages, and hence determine the mask depth with t . The mask is generated in the Evolved Mask Generation (EMG) module, which is denoted as

$$\mathcal{M}^{(t)} = \text{EMG}(\mathcal{T}^{(t)}, t). \quad (4)$$

The generated masks are hence adopted in the MIM task, and supervise the model pre-training with Eq. (2). Fig. 5 illustrates the pipeline of the proposed method. The method starts by feeding the input image into encoder $F_\theta^{(t)}$ of the current epoch t . Similarities among image patches are then used to create the hierarchical structure \mathcal{T} on the fly. Masks $\mathcal{M}^{(t)}$ are generated by performing masking at a certain depth on \mathcal{T} ,

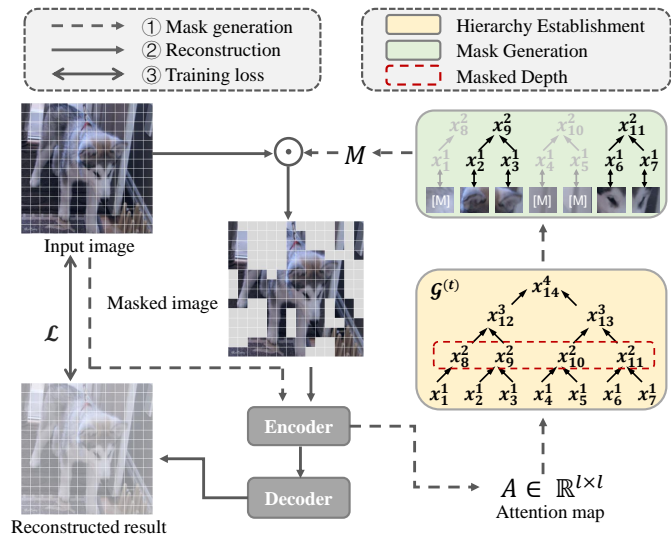


Fig. 5. The pipeline of proposed Evolved Hierarchical Masking using MAE [3] as an example. Input image is fed into the encoder extracting the attention map A to reflect the similarity between each pair of patches. The Adaptive Hierarchy Establishment module organizes leaf nodes and latent variables into the hierarchy \mathcal{T} . Based on \mathcal{T} , the Evolved Mask Generation module generates masks \mathcal{M} on a specific masking depth.

which evolves with pretraining epochs t . The masked image is adopted for MIM, which is supervised by a loss between the original and recovered images for model optimization. Our method iteratively performs the hierarchy establishment and mask generation, and finally leads to a more accurate hierarchy \mathcal{T} , and a more general pre-trained model.

The following part proceeds to present details in the Adaptive Hierarchy Establishment module and Evolved Mask Generation module, respectively. For brevity, we omit the superscript (t) denoting the training epoch in the following section.

IV. METHODOLOGY

A. Adaptive Hierarchy Establishment

Adaptive Hierarchy Establishment aims to organize image contents into a hierarchy reflecting their relationships. As the model being trained is gaining stronger visual contents modeling capability, we leverage it to compute similarity among image patches, and refer to the similarity cues for $\mathcal{T} := (V, E)$ construction. This adaptive hierarchy is dynamically constructed each time an image is input into the system. In the early stages of training, the generated hierarchy is nearly random, and the masks are similarly random. The model learns visual knowledge from these random masks and generates a more accurate hierarchy, which in turn produces more challenging masks that are appropriate to the model's current capabilities. This ensures that the model's capabilities and the difficulty of the masks are on par, fostering a kind of adversarial training mechanism that guides their co-evolution.

The hierarchy established needs to satisfy two requirements: (1) related patches are adjacent in the structure, and (2) it supports different levels of granularity. A tree structure is the most straightforward structure meets these criteria. In particular, a

binary tree is efficient and computationally lightweight, as it only requires the calculation of pair-wise similarities. Hence, we use the binary tree as a simple representation of \mathcal{T} , and formalize its establishment as an optimal problem.

Our goal is to organize the nodes into a binary tree to maximize the sum of similarity of adjacent sibling elements. Using $S(\cdot)$ as the similarity computation, the objective of \mathcal{T} can be denoted as

$$\mathcal{T} = \arg \max_{\mathcal{T}} \sum_{i=H \times W + 1}^n S(Ch_i^a), \quad (5)$$

where Ch_i^a is the set of children nodes for node x_i^a .

A series of works [14], [24] have demonstrated that transformer attention map can reflect the semantic relationship between tokens and a higher attention value represents a stronger relationship. We directly utilize the attention map from the current pre-trained model F_θ to measure the similarity between the nodes without introducing extra parameters. This strategy also ensures the accuracy of generated \mathcal{T} is on par with the modeling capacity of the trained F_θ .

Denoted as $A = \{A_{i,j}\} \in \mathbb{R}^{HW \times HW}$, the attention map produced by F_θ provides similarity cues of leaf nodes in \mathcal{T} . For instance, $A_{i,j}$ reflects the relationship between the i -th and j -th patches, which are leaf nodes x_i^1 and x_j^1 in \mathcal{T} . The similarity between two leaf nodes can be represented as

$$S(x_i^1, x_j^1) = A_{i,j}. \quad (6)$$

We hence use A as a foundation for constructing \mathcal{T} in a bottom-up way. A doesn't provide similarity cues for internal nodes at layers larger than 1. We measure their similarity by referring to their descendant leaf nodes. The similarity measurement can be computed as

$$S(x_i^a, x_j^b) = \frac{\sum_{x_k^c \in Ch_i^a} \sum_{x_m^d \in Ch_j^b} S(x_k^c, x_m^d)}{|Ch_i^a| |Ch_j^b|}, \quad (7)$$

where $|Ch|$ denotes the number of children nodes, and we define the number of children of leaf nodes as 1.

Based on the computed similarity cues of $S(\cdot)$, the training objective in Eq. (5) can be achieved by iteratively merging closely related nodes into a parent node. It starts from merging leaf nodes, and continues until the root node is generated. The initial edge set E_0 is empty and the initial node set V_0 only contains leaf nodes with cardinality $H \times W$, i.e.,

$$V_0 = \{x_1^1, x_2^1, \dots, x_{H \times W}^1\}. \quad (8)$$

At the k -th iteration, the most relevant nodes x_i^a and x_j^b are selected from V_k to form a new internal node as their parent, which can be written as

$$x_k^c = \{x_i^a, x_j^b\} = \arg \max_{(x_i^a, x_j^b)} S(x_i^a, x_j^b), x_i^a, x_j^b \in V_k. \quad (9)$$

The hierarchical level of x_k^c is determined based on the level of its child nodes, which is

$$c = \max(a, b) + 1. \quad (10)$$

The new node x_k^c and new edges are added to \mathcal{T} . Meanwhile, its children nodes are removed from \mathcal{T} , i.e.,

$$E_{k+1} = E_k \cup \{e_{a,k}, e_{b,k}\}, \quad (11)$$

$$Pa_i^a = Pa_j^b = x_k^c, Ch_k^c = \{x_i^a, x_j^b\}, \quad (12)$$

$$V_{k+1} = (V_k \setminus \{x_i^a, x_j^b\}) \cup x_k^c. \quad (13)$$

The above procedure continues for $H \times W - 1$ iterations and ends with a hierarchical structure \mathcal{T} . Note that, Eq. (7) refers to leaf nodes to compute the similarity between internal nodes. We hence do not need to compute the feature representation of each internal node. Storing node similarities into a matrix avoids recursive similarity calculation as in Eq.(7), thus accelerates the computation. There are other ways to accelerate the \mathcal{T} generation. Since the iterations are repeated for a constant $H \times W - 1$ times for every image, this procedure can be efficiently batch parallelized. The computational time complexity of hierarchy establishment is $\mathbf{O}(HW)$. HW is typically a constant 196, which constitutes only a small fraction of the total pretraining computational complexity.

B. Evolved Mask Generation

The Evolved Mask Generation module leverages the established \mathcal{T} to generate mask $\mathcal{M} \in \{0, 1\}^{H \times W}$ by referring to the current training epoch t . Note that, the training epoch t indicates the quality of established \mathcal{T} . For instance, a \mathcal{T} at the early training stage with a small t is not discriminative enough to show high-level semantics. It is hence adopted to generate simple mask patterns to reinforce the learning of low-level details in MIM. \mathcal{T} established with a larger t is hence adopted to supervise the learning of more complicated semantics.

To this end, we randomly select a portion of nodes at a specific depth h on \mathcal{T} according to t . Masks are generated by removing image patches corresponding to selected nodes. As the training progresses, h gradually moves to the upper layers of \mathcal{T} to evolve to more complicated mask patterns. h is computed based on the ratio of the current training epoch t to the total number of epochs, i.e.,

$$h = 1 + \lfloor l \times \frac{t}{total_epochs} \rfloor, \quad (14)$$

where l denotes the height of the binary tree and $total_epochs$ denotes the total number of pretraining epochs. This equation formulates the masking location h as a parameter evolves linearly from shallower to deeper layers as the pretraining progresses. It simplifies the setting of hyperparameters and works well across various tasks on different datasets in Section V.

Specifically, for nodes at the h -th layer of \mathcal{T} , we randomly select a subset of nodes, denoted as $\{x_i^h, x_j^h, \dots\}$. This selection ensures that the total count of leaf nodes corresponding to these nodes equals $HW \times r$, where r is the mask ratio. We denote the set of descendant leaf nodes of these selected nodes as V^{mask} , defined as:

$$V^{mask} = \{De_i^h, De_j^h, \dots\} \quad (15)$$

where nodes are randomly selected from the h -th depth of established tree \mathcal{T} , with the condition, that,

$$|V^{mask}| = HW \times r \quad (16)$$

If the total number of leaf nodes exceeds $HW \times r$, we truncate V^{mask} to retain only the first $HW \times r$ nodes.

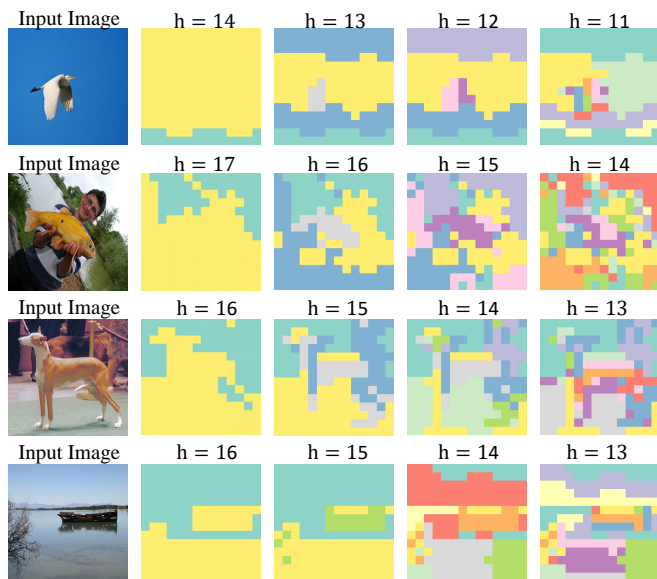


Fig. 6. Visualization of image partition at different depth h with the model after pretraining.

Finally, masks are generated by removing image patches corresponding to these selected descendant leaf nodes, that,

$$\mathcal{M}_i = \begin{cases} 0, & \text{if } x_i^1 \in V^{mask} \\ 1, & \text{otherwise.} \end{cases} \quad (17)$$

This approach ensures consistent operations across images in a batch, which makes mask generation parallelizable and enhances efficiency.

During pretraining, we elevate the masking depth to higher layers. In this way, the generated masks are adaptive to the current model parameters and evolve with the training epoch, gradually driving the network to learn from low-level details to more complex visual cue modeling. The method hence provides a progressive and more reasonable learning difficulty curve, enhancing training efficiency and yielding robust pre-trained models for a variety of downstream tasks.

C. Visualization

Established tree structure. In Fig. 6, we visualize the partitioning results of an image at varying depths h . It can be seen that as the partition height h decreases, the contents of divided partitions gradually shift from abstract concepts to specific, local parts. It's worth noting that our image partition differs from that of image segmentation: the former aims to identify relevant patches, while the latter focuses on pixel-level object division. For instance, in the first column of the figure, our method effectively divides the bird's body, a more complex region compared to the sky, into more parts. On the other hand, given the varying likelihood of the sky appearing in different parts of the image, the incorporation of positional embedding leads to a top-to-bottom division of the sky into distinct regions for masking.

Generated masks. Guided by the hierarchy structure, the generated mask contains visual hints of different abstraction levels. This guides the network to learn progressively from

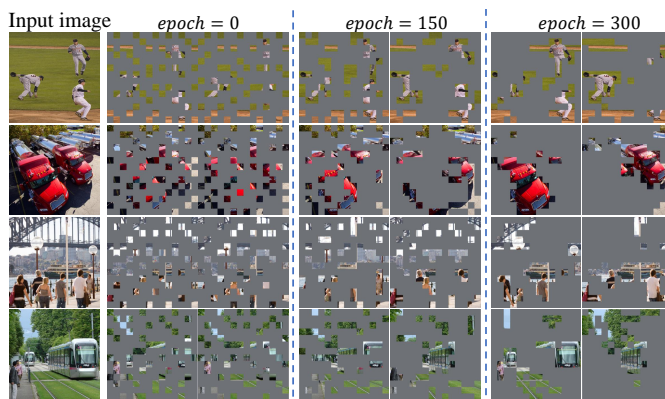


Fig. 7. Visualization of generated masks in distinct training stages. The shown masks come from using the MAE [3] method with ViT-Base [9] as the backbone, trained for 300 epochs.

simple to complex, and from specific textures to general meanings. The masks also ensure a smooth difficulty curve for the learning process. Fig. 7 visualizes the masks generated by the Evolve Hierarchy Masking in various pre-training stages. In the initial pre-training phase, the generated mask provides visual hints for most components, aiding the network in recognizing these components through texture details. This simple task helps the model to learn essential low-level visual cues and facilitates rapid convergence. As training advances, the masked patches drop more complete objects, pushing the network to build associations within and among high-level semantics. The generated masks are also more diverse, making recovery challenging while enhancing the learning of a broader range of visual knowledge. For example, in the first image of Fig. 7, the initial mask preserves visual hints for all three baseball players, grass, and playground. The network can easily infer missing pixels using adjacent similar components. After 300 training epochs, the mask largely removes some baseball players, necessitating the network to rely on the learned visual hierarchy to restore it. This enhances the capability of grasping visual cues, providing stronger initial parameters for downstream tasks.

Attention map of the pre-trained model. In Fig.2 (b), we show the size of the attended area, where these methods exhibit distinct sizes. To further explore the reasons, we visualize the attention maps of some patches in Fig. 8. It can be seen that grid-wise masking makes the pre-trained model focus on adjacent texture-similar features rather than semantically-similar ones. In the first image shown in Fig. 8, for grid-wise masking, the token on the wedding dress only focuses on the clothes patches, not the person wearing it. Since there are always nearby visible patches to provide cues, the model tends to reconstruct the masked content based on these visible patches, which makes it attend mainly to nearby resemble patches. Yet, the other masking approaches, to varying degrees, allow the network to learn more about high-level relationships. When dealing with masks covering objects with fewer visible patches, the model must learn to model the relationship between global cues to predict these objects effectively. Random masks play a role in compensating

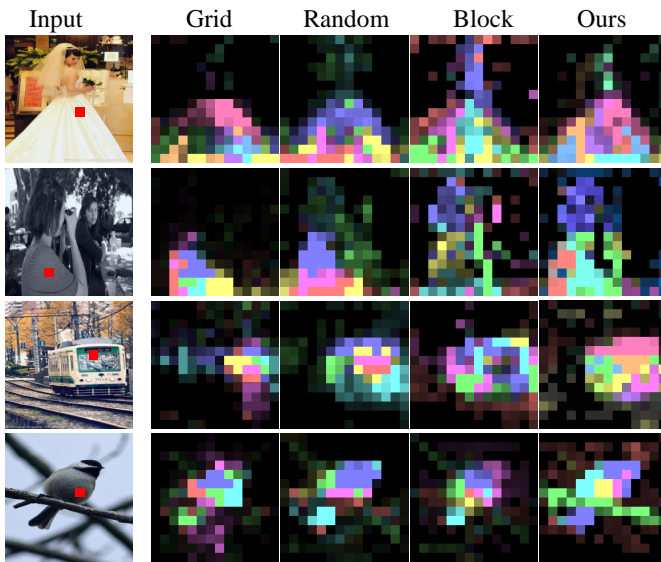


Fig. 8. **Visualization of attention map** for sampled patches under different masking methods. A red dot in the input image annotates the sampled patch. Attention heads are encoded in various colours, and the brightness indicates the attention value.

for the limitations of grid masks in terms of learning high-level semantics. It helps capture a bit more connection between objects. In comparison, block-wise masks enable the model to establish larger-scale connections between objects, though less accurately.

V. EXPERIMENTS

We evaluate our approach on seven tasks, including Image Classification, Semantic Segmentation, Detection, Object Counting (OC), Multi-Class Identification (MCI), Visual Question Answering (VQA), and Landmark Retrieval.

A. Datasets and Evaluation Metric

ImageNet-1K [10] dataset is made up of approximately 1.28M images for training and 50K images for validation with 1000 categories. Following literature [3], [4], [14], we conduct model pre-training on ImageNet-1K training images without utilizing labels. Top-1 accuracy of finetuning and linear probing are reported for classification performance.

ADE20K [11] is a challenging scene parsing benchmark, with 150 detailed semantic concepts. The training set includes 20,210 images with 150 semantic classes. The validation and test set contain 2,000 and 3,352 images respectively. In accordance with the standard evaluation method [11], we present the average mean Intersection-over-Union (mIoU) and mean Accuracy (mAcc) across all classes.

MS-COCO [49], an abbreviation for Microsoft Common Objects in Context, is a large-scale object detection, segmentation, and captioning dataset. It is comprised of more than 330K images and 1.5M object instances. It has 91 common objects found in their natural context which a 4-year-old human can easily recognize. For detection performance, we report bounding box average precision (AP-box). We report

Top-1 accuracy on MS-COCO for both object counting and multi-class identification tasks.

VCR [50], an abbreviation for Visual Commonsense Reasoning, is made up of more than 110k images for cognition-level visual understanding. Following [51], we construct the MCI task by asking the model if a certain object exists in the image with the prompt “*Question: Does obj exist in the image? Answer:*”. The model is expected to answer “*Yes/No*”, resulting in a binary classification problem. We take out the “yes” or “no” from the initial sentence the model generates and match it with the correct answer. Accuracy are reported for this task.

VQAv2 [15] is a large-scale multi-modality dataset built on the top of VQA dataset [52]. It contains approximately 1.1 Million (image, question) pairs, comprising of $\sim 900K$ abstract scene picture and $\sim 204K$ real images from COCO with approximately 13 Million associated answers. The real images dataset contains $\sim 614K$ free-form natural language questions (3 questions per image), and over 6 million free-form but concise answers (10 answers per question). Modern Language Models (MLLMs) generally produce more than one sentence. This makes it challenging to directly assess them using the standard evaluation protocol, which necessitates an exact match between predictions and ground truth. To address this, we follow [51] and modify the original evaluation protocol. We consider the first sentence generated by the MLLM as the prediction and deem it accurate if it includes the correct answer from the ground truth.

Oxford [12] and *Paris* [13] are used to evaluate our model on landmark retrieval task, referred to as *ROxf* and *RPAr*, respectively. Each dataset has 70 queries and includes 4993 and 6322 database images, respectively. Performance is measured using mean average precision (mAP). Besides, the pre-trained models are finetuned for landmark retrieval tasks on clean Google Landmarks dataset v2 (GLDv2) [53] following [54]. The clean version of GLDv2 contains 1.58M images with 81K classes.

Following GVT [51], the object counting task is constructed on MS-COCO [49] dataset by instructing the model to enumerate instances of a specific object present within an image, using the prompt “*Question: How many obj are there in the image? Answer:*”. The number in the first generated sentence is extracted as the prediction and is compared with the ground truth number. This task is treated as a classification assignment, and the performance is measured by the model’s top-1 prediction accuracy.

B. Implementation Details

We use different capacity Vision transformers as the backbones in our study, *i.e.*, ViT-S, ViT-B [9] and Swin transformer [55]. We apply the proposed masking method to three popular MIM models, *i.e.*, MAE [3], BEiT [14] and SimMIM [4]. Both pre-training and fine-tuning for classification are conducted on ImageNet-1K [10] dataset under 224×224 resolution. The images are patchified with patch size $p = 16$ for ViT and $p = 4$ for the Swin transformer. Models structure and optimization settings follow that in the corresponding

works. Pretrained models are optimized by the AdamW [56] optimizer and we set β_1 to 0.9 and β_2 to 0.95. By default, the models are pretrained for 300 epochs for a fair comparison, the mask ratio r is set to 0.75 following the original model [3].

After pre-training, the model parameters are utilized as initializations for various downstream tasks. For linear probing, we follow the standard setting in MAE [3], disabling many common regularization strategies, including mixup, cutmix, drop path and setting weight decay as zero. For semantic segmentation, we take UpperNet [57] as the framework and use the pre-trained encoder parameter to initialize the model backbone. The segmentation models are then fine-tuned on ADE20K for 160K iterations with the default setting. The input resolution is 512×512 without using multi-scale test. For object detection, we adapt the pretrained transformer to take the place of the vanilla FPN backbone [58] in Mask R-CNN [59] following [3] and fine-tune the model for 30 epochs. The ViT modules are split into 4 subsets, each performing convolutions to adjust the feature maps' sizes. This creates various scales: stride 4, 8, 16, and 32, similar to a standard ResNet. We use the AdamW [56] optimizer with a weight decay of 0.05. The standard $1 \times$ and $3 \times$ schedules are applied, where initial learning rate is 3×10^{-4} and it decays by a factor of 10 after $3/4$ and $11/12$ of finetuning epochs. For landmark retrieval, we follow DELG [54] and replace the visual backbone with the pre-trained models. These models are fine-tuned for 6M steps, optimizing with SGD and a momentum of 0.9, and a learning rate that decreases linearly to zero after reaching the desired step count.

We further evaluate the ability of the proposed pre-trained models to understand visuals based on semantic understanding and fine-grained detail perception, following GVT [51]. This assessment involves tasks involved include visual question answering, multi-class identification and object counting. They focus on evaluating high-level semantics, mid-level visual cues, and low-level details, respectively. We use different pre-trained visual encoders to represent visual contents as token sequences. These sequences are then aligned with the input space of the Large Language Model (LLM) using the Perceiver Resampler [60]. The aligned features are subsequently fed into the LLM for final prediction. The ensemble of models is fine-tuned jointly, or only the Perceiver Resampler [60] is fine-tuned while keeping the pre-trained language model fixed. We resize the input images to 224×224 and apply basic data augmentations: random resized crop and horizontal flipping. The model undergoes fine-tuning for 20K iterations using the AdamW [56] optimizer with a learning rate of 1×10^{-6} .

C. Ablation Study

Models in ablation experiments are built upon the ViT-S backbone [9] and asymmetric MAE architecture [3].

Comparison with static masking. We first investigate the performance of static masking methods and the proposed evolved method on various downstream tasks in Tab. I. It can be seen that different static masking methods exhibit distinct advantages. For example, grid-wise masking gives the network better performance with fewer pre-training epochs.

Compared with classification tasks, block-wise masking brings more performance improvements on dense prediction tasks under sufficient pre-training. The properties exhibited by random masking lie in between, including the strengths and weaknesses of the two. Meanwhile, our method combines the advantages and overcomes the disadvantages by varying masking criteria along the pre-training process, which outperforms these static methods in both performance and efficiency.

Mask ratio. The process of restoring masked areas associates the visible pixels with the removed content. This determines the visual dependencies learned by the pre-trained network, making an appropriate mask ratio crucial for pre-training. MAE [3] shows that for random masks, optimal ratios are surprisingly high. These ratios enable the network to fully grasp high-level meanings. Tab. II shows the impact of various mask ratios on random masking and the proposed method. It can be seen that the proposed method is more robust to mask ratio, outperforming random masking +1.77% and 2.24% for ImageNet classification [10] and ADE20K segmentation [11] at ratios of 0.25, respectively. The discrepancy between the highest and lowest values is merely 1.1% and 2.45% for two tasks. In comparison, the difference for random masking reaches 2.16% and 3.01% for classification and segmentation. This is because, under a mask ratio of 0.25, random masking retains too much similar content, allowing the network to restore masked patches by extending lines or textures. This issue becomes more evident in segmentation task. The proposed method takes the difficulty of pre-text task into consideration. Even at a low mask ratio, some components will be moved out as a whole, ensuring the network learns ample semantics.

Hierarchical masking. Some simpler masking methods, such as grid-like patterns for fine-grained masking and randomly large block masking for object-level masking, can also produce certain masking patterns. Tab. III compares simple masking patterns with the proposed method upon a ViT-S backbone, on ImageNet-1K classification and ADE20K segmentation tasks. We used randomly large block masking and our method with a deep masking depth to generate object-level masks. Similarly, we used grid-like masking and our method with a shallow masking depth for part-level masks. Note that we kept the depth of masking constant in our method, i.e., 20% of the height for shallow masking, and 80% for the deep one, to ensure a fair comparison. As shown in Tab. III, the proposed method achieves substantially better performance in both tasks than those simple masking methods, which demonstrates the effectiveness of hierarchical masking.

Evolved masking depth. The concern arises whether a model, after undergoing training with evolved masking depths, may forget initially acquired low-level knowledge. To investigate this issue, we compare the effects of fixed and evolved masking depths in Tab. IV. The compared masking strategies include 1) only shallow: 20% of the height as the shallow masking for low-level cues, 2) only deep: 80% of the height as the deep masking for high-level semantics, and 3) evolved: the one that evolved with training. It can be seen that, compared to using only low-level masking, the evolved method achieves substantial performance improvements in classification and segmentation tasks that relatively prioritize higher-

TABLE I

DOWNSTREAM TASKS PERFORMANCE AFTER FINE-TUNING. MODELS ARE PRE-TRAINED ON IMAGENET-1K [10] WITH DIFFERENT MASKING METHODS. WE REPORT IMAGENET-1K TOP-1 ACCURACY, ADE20K MIOU [11], AND COCO AP-BOX [49] FOR CLASSIFICATION, SEMANTIC SEGMENTATION AND OBJECT DETECTION, RESPECTIVELY. THE FIRST LINE IS THE PERFORMANCE WITH RANDOM INITIALIZATION IN FINE-TUNING. THE BEST PERFORMANCES FOR EACH TASK ARE NOTED WITH A GREY BACK COLOUR.

Epochs	Classification				Segmentation				Detection			
	grid	random	block	ours	grid	random	block	ours	grid	random	block	ours
Random ini.	71.53	71.53	71.53	71.53	21.17	21.17	21.17	21.17	19.31	19.31	19.31	19.31
100	78.28	78.11	77.85	78.29	36.65	37.60	36.02	37.82	36.11	32.70	31.78	34.62
200	78.63	79.20	78.55	79.92	36.67	38.81	36.65	39.78	36.21	34.42	33.15	36.51
400	79.11	79.42	79.29	80.41	36.78	39.43	40.67	41.37	35.19	35.21	35.08	37.45
800	79.34	79.77	79.69	80.72	36.54	39.31	41.81	42.02	34.62	38.73	38.91	39.16
Avg.	78.84	79.13	78.85	79.84	36.66	38.79	38.79	40.25	35.53	35.27	34.73	36.94

TABLE II

COMPARISON OF IMAGENET-1K CLASSIFICATION [10] AND ADE20K SEGMENTATION [11] WITH VARIOUS MASK RATIOS. WE REPORT TOP-1 ACCURACY AND MIOU RESPECTIVELY.

Mask ratio	Classification		Segmentation	
	Random	Ours	Random	Ours
0.25	77.19	78.96	36.23	38.14
0.50	79.02	79.84	38.56	39.11
0.75	79.35	80.06	39.24	40.59

TABLE III

COMPARISON BETWEEN HIERARCHICAL MASKING AND BASIC MASKING PATTERNS ON IMAGENET-1K TOP-1 ACC. AND ADE20K MIOU.

Masking Method	Classification	Segmentation
Large block	77.89	37.16
Hierarchical (only deep)	78.76	39.01
Grid-like	73.16	32.16
Hierarchical (only shallow)	79.05	39.86

TABLE IV

COMPARISON BETWEEN FIXED MASKING DEPTH AND THE EVOLVED MASKING DEPTH. MODELS, BUILT UPON ViT-S, ARE PRE-TRAINED FOR 300 EPOCHS AND EVALUATED ON OXFORD RETRIEVAL, ADE20K SEGMENTATION, AND IMAGENET-1K CLASSIFICATION.

Masking depth	Retrieval	Segmentation	Classification
Only shallow	57.11	39.86	79.05
Only deep	55.07	39.01	78.76
Evolved	57.41	41.37	80.41

level semantics, i.e., +1.51% and +1.36%, respectively. It also achieves comparable performance in retrieval task that focus on low-level semantics (57.11% vs. 57.41%). The comparison between “only deep” and “only shallow” masking shows that low-level knowledge is indeed crucial for pre-training, enabling more efficient learning and faster convergence, especially in retrieval (57.11% vs. 55.07%). This demonstrates that high-level and low-level knowledge are not mutually exclusive or conflicting but rather, low-level knowledge is foundational for learning higher-level information. The evolved masking allows for learning high-level visual semantics on the basis of low-level perception, achieving better network performance in various downstream tasks.

Training epochs. In our experiments, we selected 300 training epochs as the primary setting. Experiments in Table V investigate extending training epochs on model performance. Generally, extending pretraining duration leads to enhanced

TABLE V

EFFECT OF MORE TRAINING EPOCHS. PRETRAINED ViT-S MODELS ARE VALIDATED ON IMAGENET-1K TOP-1 ACCURACY, ADE20K MIOU [11], AND COCO AP-BOX [49].

Epochs	Classification	Segmentation	Detection
100	78.29	37.82	34.62
200	79.92	39.78	36.51
400	80.41	41.37	37.45
800	80.72	42.02	39.16
1600	81.41	42.79	40.85
2400	81.55	42.94	41.08

performance, but it also increases computational costs. Notably, the current pretraining performance has not yet reached a plateau, suggesting potential for further improvement.

D. Benchmark Performance

Applying to various MIM methods. We validate the effect of the proposed method on three popular MIM models, i.e., MAE [3], BEiT [14] and SimMIM [4] and evaluate their performance on imageNet-1K classification [10] and ADE20K segmentation [11]. For a fair comparison, the models are pre-trained using the official code for 200 epochs and fine-tuned on downstream tasks with consistent experimental settings. The results are shown in Tab. VIII. Our method notably enhances the performance of all three models, with the most significant improvement seen in the segmentation task (38.81 vs. 40.23). The performance gain varies for different MIM methods, which depends on the choice of their original masking strategies and downstream tasks. For models originally utilizing random masking, such as MAE [3] and SimMIM [4], our approach effectively boosts performance by explicitly capturing improved relationships between different parts.

Comparison with recent SSL methods. Comparison with recent SSL methods on common imageNet-1K classification settings are shown in Tab. VI. Related work ADIOS [34] is not included as it evaluates on other benchmarks. DINO [24], MoCo v3 [61] and AttnMask [6] use an extra momentum encoder as the teacher. BEiT [14] uses an extra 250M DALL-E data to pre-train dVAE tokenizer [62]. MST [5] introduces an MLP head to align the features of the teacher and student. SemMAE [8] uses a pre-trained iBOT to extract token features. Our method gets comparable performance with fewer pre-training epochs, e.g., 200 v.s. 300 in fine-tuning accuracy with

TABLE VI

COMPARISON OF POPULAR SELF-SUPERVISE LEARNING METHODS ON IMAGENET-1K [10]. EVALUATION PROTOCOLS INCLUDE TOP-1 LINEAR PROBING ACCURACY AND TOP-1 FINE-TUNING ACCURACY. ALL ENTRIES ARE ON AN IMAGE SIZE OF 224 × 224.

Method	Pre-train data	Date	Extra model	Epochs	Linear probing	Fine-tune
<i>Supervised training from scratch</i>						
ViT [9]	IN-1K w/ label	ICLR 2021	-	300	-	77.9
DeiT [31]	IN-1K w/ label	ICML 2021	-	300	-	81.8
<i>Contrastive-based SSL Pre-training</i>						
DINO [24]	IN-1K	CVPR 2021	momentum ViT	1600	74.6	82.8
MoCo v3 [61]	IN-1K	ICCV 2021	momentum ViT	300	76.5	83.2
MST [5]	IN-1K	NIPS 2021	MLP Head	100	75.0	-
AttnMask [6]	IN-1K	ECCV 2022	momentum ViT	100	76.1	-
<i>MIM SSL Pre-training</i>						
BEiT [14]	IN-1K+DALL-E	ICCV 2021	dVAE [62]	300	56.7	82.9
CAE [63]	IN-1K	arxiv 2022	-	300	64.1	83.6
MAE [3]	IN-1K	CVPR 2022	-	300	64.4	83.6
SimMIM [4]	IN-1K	CVPR 2022	-	800	56.7	83.8
SemMAE [8]	IN-1K	NIPS 2022	iBOT [33]	800	68.7	83.3
MFM [32]	IN-1K	ICLR 2023	-	300	-	83.1
LocalMIM [64]	IN-1K	CVPR 2023	-	100	-	83.3
EPM [40]	IN-1K	CVPR 2023	-	200	59.6	83.6
Ours	IN-1K	-	-	200	63.4	83.8
Ours	IN-1K	-	-	300	66.1	84.9

TABLE VII

COMPARISON OF VISUAL TOKENIZERS OF ViT-B [9] WITH DIFFERENT PRETRAINING METHODS. THE TASKS INCLUDE COMPARING VISUAL QUESTION ANSWERING, OBJECT COUNTING, AND MULTI-CLASS CLASSIFICATION ON MSCOCO [49]. ADDITIONALLY, THERE'S A MULTI-CLASS CLASSIFICATION ON VISUAL COMMONSENSE REASONING [50]. THE BEST RESULT IS MARKED WITH BOLD.

Joint Tuning	Supervised	Visual Encoder	# Pretraining Images	VQA COCO	OC COCO	MCI COCO	MCI VCR	Avg.
✗	Fully	DeiT [31]	1.28M	48.81	67.28	78.03	82.33	69.11
	Weakly	CLIP [39]	400M	51.54	64.44	78.94	83.75	69.67
	Self	DINO [24]	1.28M	49.40	66.23	73.55	79.01	67.05
		DINO v2 [65]	142M	50.56	67.91	78.63	83.88	70.25
		MAE [3]	1.28M	48.21	67.15	76.46	81.80	68.60
		Ours	1.28M	49.75	67.42	78.85	82.54	69.64
✓	Fully	DeiT [31]	1.28M	50.86	67.50	80.09	83.82	70.57
	Weakly	CLIP [39]	400M	47.87	66.71	80.98	80.94	69.13
	Self	DINO [24]	1.28M	46.28	64.35	80.17	82.90	68.43
		DINO v2 [65]	142M	49.99	65.86	79.26	82.62	69.43
		MAE [3]	1.28M	48.72	67.06	81.40	84.08	70.19
		Ours	1.28M	50.94	68.36	81.79	84.91	71.50

TABLE VIII

IMAGENET-1K [10] ACCURACY AND ADE20K [11] PERFORMANCE OF THREE POPULAR MIM MODELS BEFORE AND AFTER APPLYING THE PROPOSED MASKING WITH FINE-TUNING.

Method	Classification		Segmentation	
	top-1. acc	top-5. acc	mIoU	aAcc
MAE [3]	79.20	94.61	38.81	79.61
+ours	80.06	95.10	40.23	79.87
BEiT [14]	79.05	94.57	38.39	79.31
+ours	79.87	94.80	39.78	79.83
SimMIM [4]	78.51	94.18	38.13	79.15
+ours	79.21	94.79	39.45	79.73

state-of-art works. With significantly less computation burden, the proposed method outperforms the SOTA by 1.1%. While Our method performs modestly on linear probing, the result shows linear probing and fine-tuning results are somewhat uncorrelated. The evolved masking guides the network to learn

more linearly separable features. Since MIM methods are generally stronger non-linear feature extractors but perform unsatisfied on linear probing, we focus this work on improving fine-tune performance, i.e., a series of good initialization parameters for downstream tasks, as demonstrated in [3], [66].

Comparison with recent visual tokenizers. Recently, based on the remarkable inferring ability of LLMs, there has been a surge of research to combine LLMs and visual encoders for vision-language and even pure vision tasks. This approach, using visual encoders as the visual tokenizer for LLMs, has also become an indicator to evaluate how well visual models can extract visual information. In Tab. VII, we use GVTBench to compare our method's performance with popular visual tokenizers on various downstream tasks using the same architecture. The compared tokenizers involve supervised models (DeiT [31]), text-guided weakly supervised models ([39]), and self-supervised models (MAE [3], DINO [24], DINOv2 [65]). The evaluated tasks include Vi-

TABLE IX

SEGMENTATION PERFORMANCE ON ADE20K [11] USING THE UPERNET [57] FRAMEWORK. THE BACKBONE IS INITIALIZED WITH DEiT [31], MAE [3], CLIP [39] AND OUR PRE-TRAINED MODEL ON ViT-BASE MODEL, RESPECTIVELY.

Method	Supervised	# Pretraining Images	mIoU	aAcc	mAcc
DeiT [31]	Fully	1.28M	40.11	80.22	51.38
MAE [3]	Self	1.28M	41.24	80.62	52.77
CLIP [39]	Weakly	400M	35.54	77.65	46.56
Ours	Self	1.28M	47.26	81.97	58.92

sual Question Answering (VQA), Object Counting (OC), and Multi-Class Identification (MCI). These tasks respectively gauge the model's visual comprehension from two important perspectives: grasping semantic understanding and perceiving fine-grained visual details. VQA emphasizes the tokenizer's ability to grasp high-level semantics, while OC and MCI assess its perception of intricate visual details. During training, the language model remains frozen, while the visual tokenizer can either be frozen or optimized jointly.

The results in Tab. VII show that the proposed method narrows the gap between ImageNet-1K self-supervised techniques and large-scale pre-training. For the VQA task, the text-guided weakly supervised pre-trained model excels due to its use of extensive data for enhanced semantic understanding. CLIP surpasses supervised DeiT, self-supervised DINO, and MAE by 2.73%, 2.14%, and 3.33% without joint tuning. The proposed method built upon MAE, minimizing the gap to 1.79%. With joint tuning of visual tokenizers, this gap is further reduced to 0.60%. This is because the proposed evolved hierarchy masking method targets training at various visual modeling levels. In this way, even on small-scale pre-training data, we can extract a large number of visual dependencies, making pre-training more efficient. On the other hand, the method enhances the self-supervised approach's capability to perceive intricate details. The inherent nature of supervised and self-supervised training determines their finesse in extracting detailed information over text-guided weakly supervised. The proposed method bolsters this advantage, boosting MAE's OC COCO, MCI COCO, and MCI VCR performance by 1.30%, 0.39%, and 0.83% respectively.

Comparison on dense tasks. Tab. IX reports the segmentation performance on ADE20K [11]. The compared methods includes popular fully supervised DeiT [31], self-supervised MAE [3] and large-scale weakly supervised CLIP [39]. We take these pre-trained models to replace the backbone in the UperNet [57] framework. Compared to the commonly used fully supervised backbone and the baseline method MAE, our approach outperforms them by 7.15% and 6.02% mIoU respectively. CLIP, despite being trained weakly at a large scale, does not perform well after fine-tuning. We further visualize segmentation results of UperNet [57] in Fig. 9. Our method excels at accurately identifying segmentation categories. For example, in the second column of images, our method effectively recognizes the road, while other methods incorrectly identify it as land. It can be seen the proposed

TABLE X

LANDMARK RETRIEVAL PERFORMANCE ON OXFORD AND PARIS DATASET. WE REPORT MEAN AVERAGE PRECISION (mAP) WITH BOTH MEDIUM AND HARD EVALUATION PROTOCOLS.

Model	Supervised	Medium		Hard	
		ROxf	RPar	ROxf	RPar
DeiT [31]	Fully	56.17	72.40	31.09	54.46
CLIP [39]	Weakly	45.42	58.68	15.32	39.22
Dino [24]	Self	51.53	75.13	24.15	52.03
MAE [3]	Self	55.38	76.55	30.74	50.06
Ours	Self	59.06	79.43	33.24	58.87

hierarchical masking method effectively enhances the pre-trained model to recognize texture and semantics, making it advantageous for dense prediction tasks.

Comparison on landmark retrieval. Tab. X compares our model against the popular pretraining methods on a fine-grained visual understanding task, landmark retrieval. We present Mean Average Precision (mAP) results for both Medium and Hard splits of the Oxford and Paris image retrieval datasets. Our models are finetuned on GLDv2 for 6M steps following DELG. Dino is pre-trained on a 1.2M clean set of GLDv2. It can be seen that self-supervised and supervised methods outperform large-scale weakly supervised CLIP on fine-grained perception. The proposed approach significantly enhances the fine-grained perception ability of pre-trained models, demonstrating the effectiveness of hierarchical masking in learning texture details.

E. Discussion

This section studies the properties of masking methods and analyses the causes of their performance gap.

Training costs. We measured both the space and time consumption of the proposed masking strategy during pretraining. The results show that mask generation occupies only 1.3% of the total memory space and 8.1% of the pretraining time. Specifically, obtaining the attention map accounts for 7.9% of the total computation time, while constructing the binary tree takes up only 0.2%. In our implementation, attention map A is acquired by averaging the outputs of the initial layers, i.e., the first 6 layers in our experiments, instead of fully engaging the entire model depth. Furthermore, we share the features after the embedding layer to avoid redundant computations. The experimental results demonstrate the efficiency of the proposed method. Compared to the saved training epochs, the increased time consumption is minor.

Fine-tuning gradient. The lottery ticket hypothesis [67] demonstrates winning ticket weights tend to change by a larger amount than weights in the rest of the network, which is accompanied by larger gradients. Fig. 11 shows the fine-tuning gradients for different initialized models. It can be seen that grid-wise masks better help model convergence in shallow layers. And block-wise masking is more helpful for training deep layers, promoting the model learning high-level semantic relationships. The proposed evolved masking method facilitates both deep and shallow network layers.

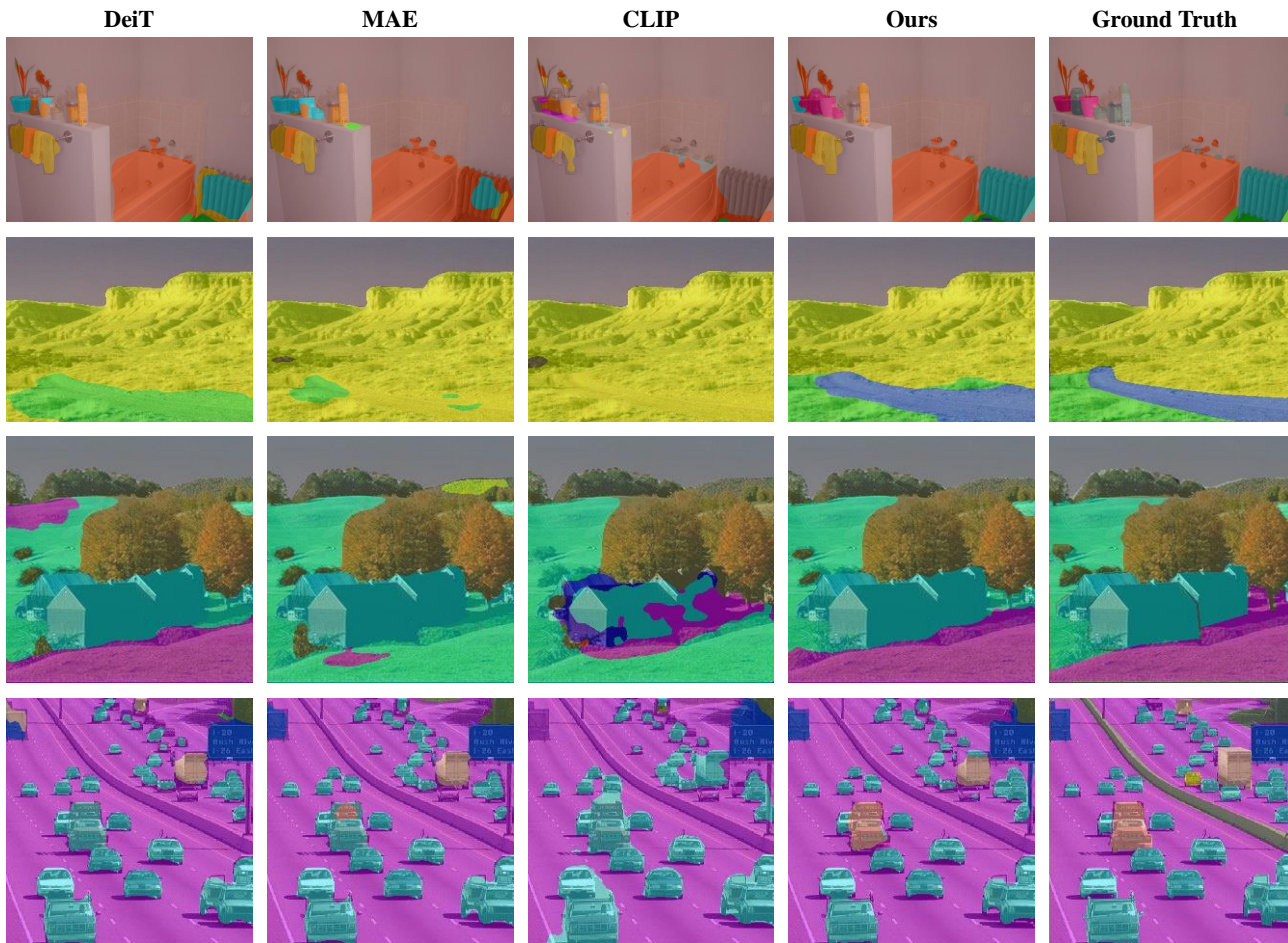


Fig. 9. Visualized segmentation results on ADE20K [11] with DeiT [31], MAE [3], CLIP [39], and the proposed method pretrained, respectively. Different colors denote different objects. Best viewed in color pdf.

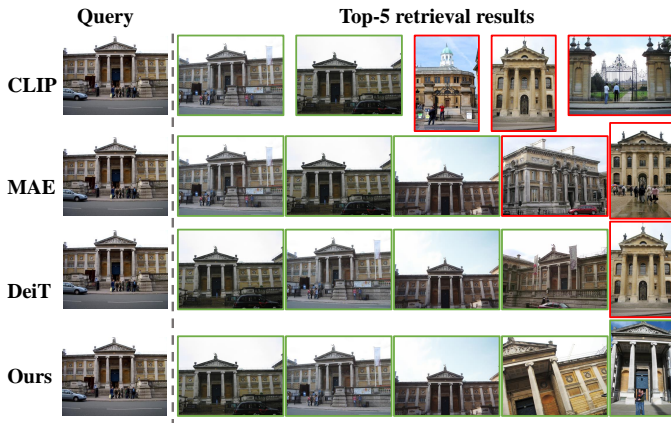


Fig. 10. Visualization of top-5 retrieval results on Oxford dataset [12]. Green and red boxes indicate correct and wrong predictions, respectively.

To summarize, the visual hints shared in visible and masked patches determine the level of abstraction at which visual knowledge is learned. The masks where visible and masked patches with similar hints make the pre-trained model learn more low-level texture knowledge and better facilitate model convergence in shallow layers during fine-tuning. Masking

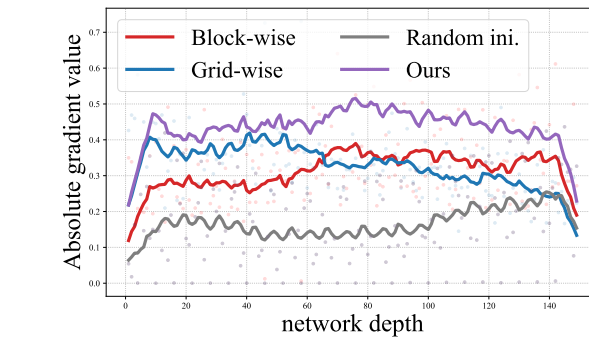


Fig. 11. Fine-tuning gradient value against network depth on imageNet-1K classification [10]. The data are fitted using exponential moving averages for better visualization. The gray line represents the model with random initialization of parameters, without pre-training.

visual cues at a higher abstraction level prompts the model to pay more attention to global information and learn the connection between objects, which benefits more on the convergence of deep layers and dense prediction tasks. The proposed method allows masks to adapt with model training, facilitating effective learning across various visual modeling

levels, and providing stronger initialization parameters for pretraining-agnostic downstream tasks.

VI. CONCLUSION

This paper proposes to enhance existing MIM methods by designing an evolved hierarchy masking strategy. We found that the design of masks directly determines the visual knowledge learned by the network in each iteration. The proposed method leverages the current training network to analyze images, capturing relationships among potential visual cues without requiring an additional network. Subsequently, it selectively masks patches corresponding to some chosen cues. Through this technique, we can extract abundant visual dependencies from a limited pre-training dataset, making pre-training more focused and effective. Experimental results demonstrate that the proposed method effectively improves the efficiency of pre-training, and has achieved performance improvements in seven downstream tasks. When combined with Large Language Models, the proposed method not only underscores the advantages of self-supervised pre-training but also narrows the gap compared to extensive pre-training for semantic tasks. We hope that the proposed method can provide timely insights for a cost-effective pathway towards self-supervised large-scale pre-training.

ACKNOWLEDGEMENTS

This work is supported in part by Grant No. 2023-JCJQ-LA-001-088, in part by Natural Science Foundation of China under Grant No. U20B2052, 61936011, in part by the Okawa Foundation Research Award, in part by the Ant Group Research Fund, and in part by the Kunpeng&Ascend Center of Excellence, Peking University.


REFERENCES

- [1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [2] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [3] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.
- [4] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, and H. Hu, "SimMIM: A simple framework for masked image modeling," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 9653–9663.
- [5] Z. Li, Z. Chen, F. Yang, W. Li, Y. Zhu, C. Zhao, R. Deng, L. Wu, R. Zhao, M. Tang *et al.*, "Mst: Masked self-supervised transformer for visual representation," *Advances in Neural Information Processing Systems*, vol. 34, pp. 13 165–13 176, 2021.
- [6] I. Kakogeorgiou, S. Gidaris, B. Psomas, Y. Avrithis, A. Bursuc, K. Karantzas, and N. Komodakis, "What to hide from your students: Attention-guided masked image modeling," *arXiv preprint arXiv:2203.12719*, 2022.
- [7] M. Assran, Q. Duval, I. Misra, P. Bojanowski, P. Vincent, M. Rabat, Y. LeCun, and N. Ballas, "Self-supervised learning from images with a joint-embedding predictive architecture," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 15 619–15 629.
- [8] G. Li, H. Zheng, D. Liu, C. Wang, B. Su, and C. Zheng, "Semmae: Semantic-guided masking for learning masked autoencoders," *arXiv preprint arXiv:2206.10207*, 2022.
- [9] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [11] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, "Scene parsing through ade20k dataset," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 633–641.
- [12] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman, "Object retrieval with large vocabularies and fast spatial matching," in *2007 IEEE conference on computer vision and pattern recognition*. IEEE, 2007, pp. 1–8.
- [13] —, "Lost in quantization: Improving particular object retrieval in large scale image databases," in *2008 IEEE conference on computer vision and pattern recognition*. IEEE, 2008, pp. 1–8.
- [14] H. Bao, L. Dong, and F. Wei, "Beit: Bert pre-training of image transformers," *arXiv preprint arXiv:2106.08254*, 2021.
- [15] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh, "Making the v in vqa matter: Elevating the role of image understanding in visual question answering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 6904–6913.
- [16] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06)*, vol. 2. IEEE, 2006, pp. 1735–1742.
- [17] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [18] X. Wang, R. Zhang, C. Shen, T. Kong, and L. Li, "Dense contrastive learning for self-supervised visual pre-training," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3024–3033.
- [19] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," *arXiv preprint arXiv:2003.04297*, 2020.
- [20] J. Zhang and K. Ma, "Rethinking the augmentation module in contrastive learning: Learning hierarchical augmentation invariance with expanded views," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 650–16 659.
- [21] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 9729–9738.
- [22] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9912–9924, 2020.
- [23] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar *et al.*, "Bootstrap your own latent—a new approach to self-supervised learning," *Advances in neural information processing systems*, vol. 33, pp. 21 271–21 284, 2020.
- [24] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9650–9660.
- [25] S. Gidaris, P. Singh, and N. Komodakis, "Unsupervised representation learning by predicting image rotations," *arXiv preprint arXiv:1803.07728*, 2018.
- [26] M. Noroozi and P. Favaro, "Unsupervised learning of visual representations by solving jigsaw puzzles," in *European conference on computer vision*. Springer, 2016, pp. 69–84.
- [27] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proceedings of the 25th international conference on Machine learning*, 2008, pp. 1096–1103.
- [28] M. Chen, A. Radford, R. Child, J. Wu, H. Jun, D. Luan, and I. Sutskever, "Generative pretraining from pixels," in *International conference on machine learning*. PMLR, 2020, pp. 1691–1703.
- [29] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context encoders: Feature learning by inpainting," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2536–2544.
- [30] R. Zhang, P. Isola, and A. A. Efros, "Colorful image colorization," in *European conference on computer vision*. Springer, 2016, pp. 649–666.

- [31] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *International Conference on Machine Learning*, PMLR, 2021, pp. 10 347–10 357.
- [32] J. Xie, W. Li, X. Zhan, Z. Liu, Y. S. Ong, and C. C. Loy, "Masked frequency modeling for self-supervised visual pre-training," *arXiv preprint arXiv:2206.07706*, 2022.
- [33] J. Zhou, C. Wei, H. Wang, W. Shen, C. Xie, A. Yuille, and T. Kong, "ibot: Image bert pre-training with online tokenizer," *arXiv preprint arXiv:2111.07832*, 2021.
- [34] Y. Shi, N. Siddharth, P. Torr, and A. R. Kosiorek, "Adversarial masking for self-supervised learning," in *International Conference on Machine Learning*. PMLR, 2022, pp. 20 026–20 040.
- [35] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. E. Gonzalez, and I. Stoica, "Judging llm-as-a-judge with mt-bench and chatbot arena," 2023.
- [36] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," *arXiv preprint arXiv:2304.08485*, 2023.
- [37] D. Zhu, J. Chen, X. Shen, X. Li, and M. Elhoseiny, "Minigt-4: Enhancing vision-language understanding with advanced large language models," *arXiv preprint arXiv:2304.10592*, 2023.
- [38] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," *arXiv preprint arXiv:2301.12597*, 2023.
- [39] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [40] Z. Feng and S. Zhang, "Evolved part masking for self-supervised learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 10 386–10 395.
- [41] A. Anandkumar, D. Hsu, A. Javanmard, and S. Kakade, "Learning linear bayesian networks with latent variables," in *International Conference on Machine Learning*. PMLR, 2013, pp. 249–257.
- [42] B. Huang, C. J. H. Low, F. Xie, C. Glymour, and K. Zhang, "Latent hierarchical causal structure discovery with rank constraints," *Advances in Neural Information Processing Systems*, vol. 35, pp. 5549–5561, 2022.
- [43] L. Kong, M. Q. Ma, G. Chen, E. P. Xing, Y. Chi, L.-P. Morency, and K. Zhang, "Understanding masked autoencoders via hierarchical latent variable models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7918–7928.
- [44] G. Hinton, "How to represent part-whole hierarchies in a neural network," *Neural Computation*, pp. 1–40, 2022.
- [45] S. Sabour, N. Frosst, and G. E. Hinton, "Dynamic routing between capsules," *Advances in neural information processing systems*, vol. 30, 2017.
- [46] L. Maaløe, M. Fraccaro, V. Liévin, and O. Winther, "Biva: A very deep hierarchy of latent variables for generative modeling," *Advances in neural information processing systems*, vol. 32, 2019.
- [47] A. Vahdat and J. Kautz, "Nvae: A deep hierarchical variational autoencoder," *Advances in neural information processing systems*, vol. 33, pp. 19 667–19 679, 2020.
- [48] S. Zhao, J. Song, and S. Ermon, "Learning hierarchical features from generative models," *arXiv preprint arXiv:1702.08396*, 2017.
- [49] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.
- [50] R. Zellers, Y. Bisk, A. Farhadi, and Y. Choi, "From recognition to cognition: Visual commonsense reasoning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [51] G. Wang, Y. Ge, X. Ding, M. Kankanhalli, and Y. Shan, "What makes for good visual tokenizers for large language models?" *arXiv preprint arXiv:2305.12223*, 2023.
- [52] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh, "Vqa: Visual question answering," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2425–2433.
- [53] T. Weyand, A. Araujo, B. Cao, and J. Sim, "Google Landmarks Dataset v2 - A Large-Scale Benchmark for Instance-Level Recognition and Retrieval," in *Proc. CVPR*, 2020.
- [54] B. Cao, A. Araujo, and J. Sim, "Unifying deep local and global features for image search," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*. Springer, 2020, pp. 726–743.
- [55] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.
- [56] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [57] T. Xiao, Y. Liu, B. Zhou, Y. Jiang, and J. Sun, "Unified perceptual parsing for scene understanding," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 418–434.
- [58] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [59] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [60] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira, "Perceiver: General perception with iterative attention," in *International conference on machine learning*. PMLR, 2021, pp. 4651–4664.
- [61] X. Chen, S. Xie, and K. He, "An empirical study of training self-supervised vision transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9640–9649.
- [62] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, "Zero-shot text-to-image generation," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8821–8831.
- [63] X. Chen, M. Ding, X. Wang, Y. Xin, S. Mo, Y. Wang, S. Han, P. Luo, G. Zeng, and J. Wang, "Context autoencoder for self-supervised representation learning," *arXiv preprint arXiv:2202.03026*, 2022.
- [64] H. Wang, Y. Tang, Y. Wang, J. Guo, Z.-H. Deng, and K. Han, "Masked image modeling with local multi-scale reconstruction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 2122–2131.
- [65] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby *et al.*, "Dinov2: Learning robust visual features without supervision," *arXiv preprint arXiv:2304.07193*, 2023.
- [66] X. Chen and K. He, "Exploring simple siamese representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 750–15 758.
- [67] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," *arXiv preprint arXiv:1803.03635*, 2018.



Zhanzhou Feng is a PhD student at School of Computer Science, Peking University, Beijing, China. He received the B.S. degree at College of Intelligence and Computing, Tianjin University in 2021. His current research interests are computer vision and deep learning, with focus on efficient neural architecture design and self-supervised learning.



Shiliang Zhang received the Ph.D. degree in computer science from the Institute of Computing Technology, Chinese Academy of Sciences. He was a Post-Doctoral Scientist with NEC Laboratories America and a Post-Doctoral Research Fellow with The University of Texas at San Antonio. He is currently an Associate Professor with Tenure with the Department of Computer Science, School of Electronic Engineering and Computer Science, Peking University.

His research interests include large-scale image retrieval and computer vision. He has authored or co-authored over 80 papers in journals and conferences, including IJCV, IEEE Trans. on PAMI, IEEE Trans. on Image Processing, IEEE Trans. on NNLS, IEEE Trans. on Multimedia, ACM Multimedia, ICCV, CVPR, ECCV, AAAI, IJCAI, etc. He was a recipient of the Outstanding Doctoral Dissertation Awards from the Chinese Academy of Sciences and Chinese Computer Federation, the President Scholarship from the Chinese Academy of Sciences, the NEC Laboratories America Spot Recognition Award, the NVidia Pioneering Research Award, and the Microsoft Research Fellowship. He was a recipient of the Top 10% Paper Award at the IEEE MMSP 2011. He served as the Associate Editor (AE) of Computer Vision and Image Understanding (CVIU) and IET Computer Vision, Guest Editor of ACM TOMM, and Area Chair or Senior Program Committee of CVPR, AAAI, ICPR, and VCIP. His research is supported by the The National Key Research and Development Program of China, Natural Science Foundation of China, Beijing Natural Science Foundation, and Microsoft Research, etc.