

Wavelet-Driven Masked Image Modeling: A Path to Efficient Visual Representation

Wenzhao Xiang^{1,2,3}, Chang Liu⁴, Hongyang Yu^{2*}, Xilin Chen^{1,3}

¹Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences

²Peng Cheng Laboratory

³University of Chinese Academy of Sciences

⁴Department of Electronic Engineering, Shanghai Jiao Tong University

xiangwenzhao22@mails.ucas.ac.cn, sunrise6513@sjtu.edu.cn, yuhy01@pcl.ac.cn, xlchen@ict.ac.cn

Abstract

Masked Image Modeling (MIM) has garnered significant attention in self-supervised learning, thanks to its impressive capacity to learn scalable visual representations tailored for downstream tasks. However, images inherently contain abundant redundant information, leading the pixel-based MIM reconstruction process to focus excessively on finer details such as textures, thus prolonging training times unnecessarily. Addressing this challenge requires a shift towards a compact representation of features during MIM reconstruction. Frequency domain analysis provides a promising avenue for achieving compact image feature representation. In contrast to the commonly used Fourier transform, wavelet transform not only offers frequency information but also preserves spatial characteristics and multi-level features of the image. Additionally, the multi-level decomposition process of wavelet transformation aligns well with the hierarchical architecture of modern neural networks. In this study, we leverage wavelet transform as a tool for efficient representation learning to expedite the training process of MIM. Specifically, we conduct multi-level decomposition of images using wavelet transform, utilizing wavelet coefficients from different levels to construct distinct reconstruction targets representing various frequencies and scales. These reconstruction targets are then integrated into the MIM process, with adjustable weights assigned to prioritize the most crucial information. Extensive experiments demonstrate that our method achieves comparable or superior performance across various downstream tasks while exhibiting higher training efficiency.

Introduction

Driven by the success of masked language modeling (MLM) in natural language processing (NLP) (Kenton and Toutanova 2019; Liu et al. 2019; Yang et al. 2019), masked image modeling (MIM) (Bao et al. 2021; He et al. 2022; Xie et al. 2022b) has emerged as a leading approach in the realm of self-supervised visual representation learning. MIM helps the model learn scalable and rich representations for various downstream tasks by masking a portion of the input data and then predicting the missing data based on the visible parts. As pixel-based MIM methods (He et al. 2022; Xie et al.

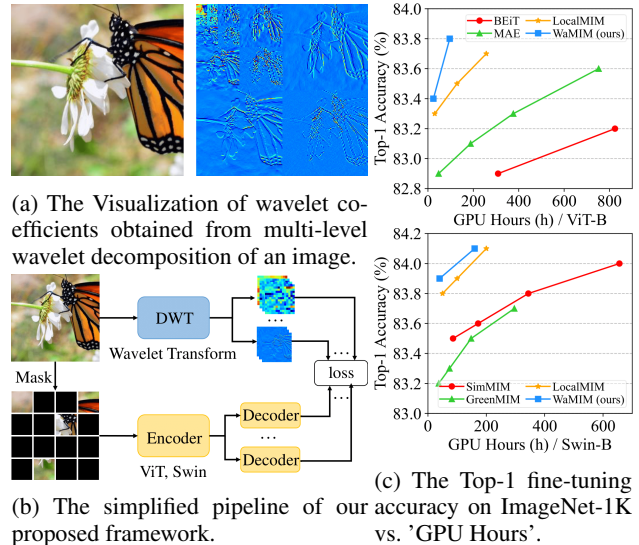


Figure 1: An overview of the proposed WaMIM.

2022b) have developed, researchers have begun directly employing raw pixel values as reconstruction targets, showcasing a simple and effective approach.

While using raw pixel values as reconstruction targets preserves the original information in MIM, the raw pixels of an image often contain a substantial amount of redundant information that isn't particularly beneficial for representation learning. Some studies (Liu et al. 2023c,d) suggest that directly predicting raw pixel values can make pre-training overly sensitive to fine-grained details of images, such as textures. Consequently, this can render the self-supervised pre-training process more challenging and unstable, requiring prolonged training to acquire rich and effective representations for downstream tasks. This necessitates a discriminative and compact feature representation for the reconstruction targets in the MIM process.

Some studies (Xie et al. 2022a; Liu et al. 2023c,d) have tackled these issues by shifting their focus to the frequency domain. They achieve this by partitioning frequency bands and utilizing more specific low-frequency or high-frequency information to construct compact reconstruction targets. These methods commonly employ the Fourier transform

*Corresponding author.

for frequency domain analysis. However, the Fourier transform only provides global frequency characteristics of signals while completely disregarding their local spatial properties, leading to the loss of valuable information. Compared to the Fourier transform, the wavelet transform preserves the spatial properties of an image while providing its frequency characteristics (Arts and van den Broek 2022). Additionally, the multi-level decomposition mechanism helps extract multi-scale features of the image, which are highly beneficial in enhancing both the speed and effectiveness of MIM (Wang et al. 2023; Ren et al. 2023).

In this paper, we introduce an efficient MIM training framework, leveraging wavelet transform to build the reconstruction targets from both frequency and multi-scale perspectives. We first perform wavelet multi-level decomposition on the input image, obtaining multi-level wavelet coefficients representing different frequency bands and scales. Then, a subset of these wavelet coefficients is selected to construct the multi-level reconstruction targets. Given that wavelet coefficients inherently act as feature descriptors reflecting features of different frequencies and scales within the image (Said et al. 2016), as depicted in Figure 1a, our framework directly employs these coefficients as multi-level reconstruction targets. The output features from different encoder layers are decoded to predict the targets, with the shallow layers corresponding to high-frequency information and deep layers to low-frequency information. This strategy is based on the observation that shallow layers of neural networks tend to capture low-level details of images, while deeper layers tend to capture high-level semantics (Park and Kim 2022; Gao et al. 2022). The loss weights of different reconstruction targets can be adjusted to prioritize more crucial features during pre-training, thereby accelerating the representation learning process. We denote the proposed method as "Wavelet-Driven Masked Image Modeling" (WaMIM), as illustrated in Figure 1.

Overall, our contributions can be summarized as follows:

- Our approach introduces a framework to accelerate and enhance MIM pre-training by incorporating multi-level reconstruction targets, generated with the guidance of wavelet transform in both spatial and frequency domains. To the best of our knowledge, we are the first to introduce wavelet guidance into the MIM pre-training.
- We propose to directly utilize the wavelet coefficients as the multi-level reconstruction targets, which can be effortlessly integrated into most existing MIM frameworks.
- Extensive experiments demonstrate that our method can accelerate and enhance the MIM methods, as shown in Figure 1c and Table 1. For example, WaMIM achieves a fine-tuning accuracy of 82.0%/83.8% on ViT-S/B with only 45/96 GPU hours, reducing the computational cost to just 47%/13% of MAE while demonstrating similar or slightly improved performance (1.1%/0.2% increase). Similar experimental results can be observed on the Swin Transformer and other downstream tasks.

Related Work

Masked Image Modeling. Building on the success of masked language modeling (MLM), particularly with mod-

els like BERT (Kenton and Toutanova 2019), masked image modeling (MIM) has emerged as a parallel approach in the visual domain, learning robust visual representations by predicting masked image regions. BEiT (Bao et al. 2021) is one of the pioneering methods in this space, utilizing an offline tokenizer, VQ-VAE (Van Den Oord, Vinyals et al. 2017), to convert images into discrete visual tokens and establish a patch-level dictionary, to recover the token IDs for the masked patches. iBOT (Zhou et al. 2022) takes a different approach by employing an online tokenizer, which generates targets for the encoder in a self-distillation framework for pre-training. Other methods, such as MAE (He et al. 2022) and SimMIM (Xie et al. 2022b), shift the focus to directly reconstructing raw pixel values, with MAE concentrating on visible patches and SimMIM on all patches. MaskFeat (Wei et al. 2022), on the other hand, avoids predicting raw pixel values by using low-level features like Histograms of Oriented Gradients (HOG) (Dalal and Triggs 2005) as reconstruction targets. In addition to these, some approaches (Wang et al. 2023; Ren et al. 2023) explore the use of deep features, frequencies, or introduce novel pretext tasks like reconstructing corrupted images (Fang et al. 2022) or denoising images (You et al. 2024; Xiang et al. 2024). More recent research (Wang et al. 2023; Ren et al. 2023) has proposed integrating multi-level supervision mechanisms into MIM, leading to significant improvements in both the speed and effectiveness of representation learning.

Frequency Domain Analysis. Frequency domain analysis has been extensively utilized across numerous computer vision tasks, including representation learning (Xu et al. 2020; Li et al. 2023; Liu et al. 2023a; Zhu and Soricut 2024), image generation (Jiang et al. 2021; Phung, Dao, and Tran 2023), and image super-resolution (Fuoli, Van Gool, and Timofte 2021; Liu, Cheng, and Tan 2023). For a long time, frequency analysis, encompassing Fourier analysis and wavelet analysis, has served as a pivotal tool in handling signals due to its inherent capability to segregate information across various scales. The efficacy of frequency analysis in expediting the training process of deep neural networks has been corroborated, as demonstrated in previous works such as (Yao et al. 2022; Park and Johnson 2023). The efficiency of training MIM models is also dependent on a more effective and dense representation of image information. Besides, some findings indicate that shallow layers in MAE encompass notably more high-frequency components compared to deep layers, which primarily capture low-level details such as textures (Liu et al. 2023d). This insight suggests the potential of generating targets based on information in the frequency domain. Several approaches (Xie et al. 2022a; Liu et al. 2023b,c) leverage frequency analysis to extract crucial features for reconstruction targets.

Proposed Method

Preliminaries

Masked Image Modeling. In the MIM approach, models are tasked with predicting information from images with certain regions masked. To achieve this objective, the overarching framework comprises three main components: image

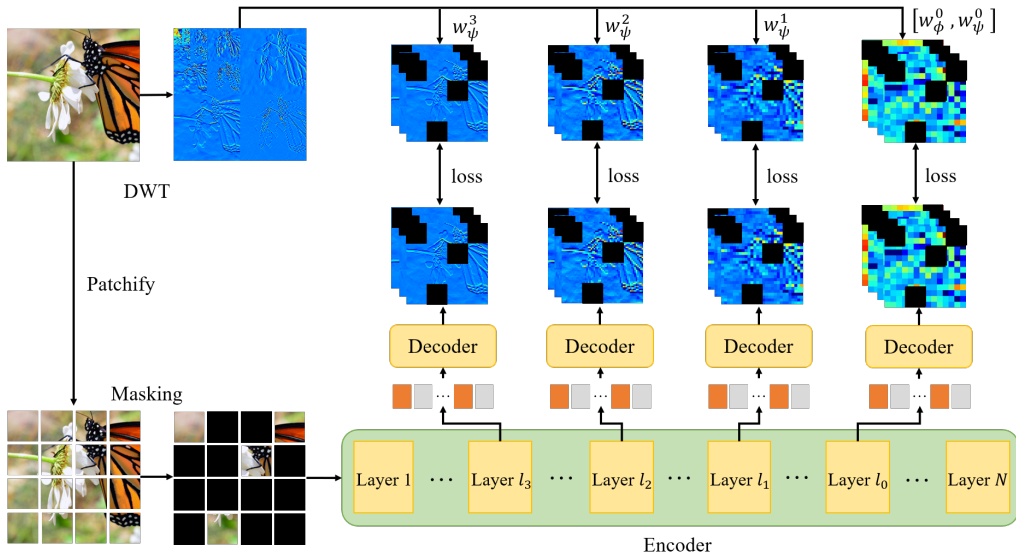


Figure 2: The training framework of WaMIM. The images are first patchified and masked. They are then passed through an encoder to extract features from various layers. The decoder selects features from specific layers to reconstruct the targets, which are generated using discrete wavelet transform (DWT) to decompose the images into different scales. The wavelet coefficient w_{ψ}^{K-k} serves as the reconstruction target for features extracted from Layer l_k . The larger k corresponds to deeper layers in the encoder, aim at reconstructing higher frequency wavelet coefficients. K represents the total number of layers selected for reconstruction. For $k = K$, the coefficient of the scaling function is concatenated with that of the wavelet function.

masking, an auto-encoder model, and a target generator.

Formally, let $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$ denote the input image, where H , W , and C represent the height, width, and number of channels of the image respectively. The input images undergo masking with f_m , denoted as $\hat{\mathbf{x}} = f_m(\mathbf{x})$. The masking operation reorganizes the pixels into $N = HW/p^2$ non-overlapping patches, where p represents the patch size, and masks the patches with a random mask $m \in \{0, 1\}^N$, where $m_i = 1$ indicates the patch is masked.

The corrupted image is then fed into an auto-encoder f_{AE} to predict the target information of the original image. Typically, the auto-encoder comprises an encoder f_{EN} and a decoder f_{DE} . Assuming the encoder has L layers, the output of the encoder contains features from different scales, i.e., $[z^1, z^2, \dots, z^L] = f_{EN}(\hat{\mathbf{x}})$, and the output of the last layer of the encoder $z^L = f_{EN}^L(\hat{\mathbf{x}})$ is decoded as $\hat{\mathbf{y}}^L = f_{DE}(z^L)$, representing the prediction for the target \mathbf{y} . This concept is further extended in (Wang et al. 2023), which suggests that local multi-scale reconstruction offers a more efficient approach. Consequently, a more comprehensive objective for MIM involves reconstructing information from multiple layers. For each latent code z^l , where $l \in [1, 2, \dots, L]$, the information is reconstructed as $\hat{\mathbf{y}}^l = f_{DE}^l(z^l)$, with \mathbf{y}^l representing the target. The reconstruction target is crucial in MIM and is generated by a target generator function f_t , such that $\mathbf{y}^l = f_t^l(\mathbf{x})$. Combining these elements yields the general formulation for MIM as

$$\mathcal{L} = \sum_{l=1}^L c^l \cdot \mathcal{D}(f_{AE}^l(f_m(\mathbf{x})), f_t^l(\mathbf{x})), \quad (1)$$

in which, \mathcal{D} measures the distance (e.g., L1 or L2 distance) between the prediction and the target, $f_{AE}^l = f_{DE}^l \circ f_{EN}^l$, c^l is the weight for the reconstruction loss of each layer. It should be noted that the distance measurement $\mathcal{D}(\cdot, \cdot)$ implicitly contains the operation of the masking operation to accelerate calculation. From this perspective, MAE (He et al. 2022) is a special case of MIM with $c^l = 1$ for $l = L$, $c^l = 0$ for $l \neq L$ and $f_t^l(\mathbf{x}) = \mathbf{x}$.

Wavelet Transform. Different from Fourier analysis, which examines signals solely in the frequency domain, wavelet analysis strikes a balance between time-domain and frequency-domain analysis. The wavelet transform decomposes signals using a dilated and translated version of a scaling function $\phi(t)$ and a wavelet function $\psi(t)$, both of which are defined as signals with zero mean and finite energy. In this paper, our focus lies on 2D discrete wavelet transform (DWT). Here, the scaling function is expanded as $\phi(t_m, t_n)$, and the wavelet function is expanded into three two-dimensional functions: $\psi^H(t_m, t_n)$ (measuring variations along columns), $\psi^V(t_m, t_n)$ (measuring variations along rows), and $\psi^D(t_m, t_n)$ (measuring variations along diagonals). Utilizing these scaling and wavelet functions, we define the scaled and translated basis functions as

$$\begin{aligned} \phi_{j,m,n}(t_m, t_n) &= 2^{\frac{j}{2}} \phi(2^j t_m - m, 2^j t_n - n), \quad (2) \\ \psi_{j,m,n}^i(t_m, t_n) &= 2^{\frac{j}{2}} \psi^i(2^j t_m - m, 2^j t_n - n), \quad i = H, V, D, \quad (3) \end{aligned}$$

where index i identifies the direction of the wavelet function.

The space spanned by $\phi_{j,m,n}$ is nested, with j increasing, resulting in the entire space being defined within $L^2(\mathbb{R})$. The

space spanned by $\psi_{j,m,n}^i$ is orthogonal to the space spanned by $\phi_{j,m,n}$ at the same j , ensuring that the combination of $\phi_{j_0,m,n}, \psi_{j_0,m,n}^i, \dots, \psi_{j_\infty,m,n}^i$ forms the basis to represent two-dimensional signals. For an image \mathbf{x} with size $M \times N$, we typically set $j_0 = 0$ and choose $N = M = 2^J$, so that $j = 0, 1, \dots, J$ and $m = n = 0, 1, \dots, 2^j - 1$. The wavelet transform can be expressed as

$$[\mathbf{w}_\phi^0; \mathbf{w}_\psi^0, \dots, \mathbf{w}_\psi^J] = f_w(\mathbf{x}), \quad (4)$$

where f_w is the forward function of DWT. Detailed expression is shown as

$$\begin{aligned} \mathbf{w}_\phi^0[m, n] &= \frac{1}{\sqrt{MN}} \sum_{t_m=0}^{M-1} \sum_{t_n=0}^{N-1} p(t_m, t_n) \phi_{0,m,n}(t_m, t_n), \\ \mathbf{w}_\psi^j[i, m, n] &= \frac{1}{\sqrt{MN}} \sum_{t_m=0}^{M-1} \sum_{t_n=0}^{N-1} p(t_m, t_n) \psi_{j,m,n}^i(t_m, t_n), \end{aligned} \quad (5)$$

where $i = \{H, V, D\}$, $p(t_m, t_n)$ is the pixel value function of \mathbf{x} . The expression provided above is a common representation of the DWT. In practice, DWT is often implemented using Multi-Resolution Analysis, which facilitates a fast algorithm and eliminates redundant information.

Wavelet-Driven Reconstruction Target

Creating reconstruction targets is a critical component of the MIM method. Effective reconstruction targets not only aid the model in acquiring superior representations for downstream tasks but also enhance pre-training efficiency and conserve computational resources. Following the works of MAE and SimMIM, many methods have employed raw or normalized pixel values as reconstruction targets. However, raw pixel values contain a considerable amount of redundant information. Directly employing them as reconstruction targets may result in the reconstruction task being overly sensitive to the high-frequency components of the image, which could be detrimental to representation learning.

Using frequency analysis tools can address the mentioned issue by selectively processing various frequency bands of the reconstruction targets. The Fourier transform is the most common tool for frequency domain analysis as in (Xie et al. 2022a; Liu et al. 2023b,c). However, it only offers the global frequency response of the signal. Relying solely on the Fourier transform may result in the loss of local spatial characteristics and multi-scale information in the input image, which are also critical factors in enhancing representation learning (Wang et al. 2023; Ren et al. 2023).

In our method, we utilize the wavelet transform to guide the construction of reconstruction targets. In contrast to the Fourier transform, the wavelet transform, as shown in Equation 4, maintains both spatial and frequency characteristics of an image, with its multi-level decomposition mechanism facilitating the extraction of multi-scale features of the image. As the wavelet coefficients are commonly used as the feature descriptor of input signals, for simplicity, we directly set the wavelet coefficients from DWT as the targets of the generator f_t . As j in Equation 5 increases, the coefficients refer to information with higher frequency, which

are related to low-level features, *i.e.*, features extracted from the shallow layer of the encoder. Therefore, we select K ($K = J + 1, K \leq L$) sets of features extracted from layers $\{l_1, l_2, \dots, l_K\} \subseteq \{1, 2, \dots, L\}$. Then, we set the target for the features encoded from the l_k -th layer ($k \in \{1, \dots, K\}$) as the wavelet coefficients of $j = K - k$, which is shown as

$$f_t^{l_k}(\mathbf{x}) = \begin{cases} \mathbf{w}_\psi^{K-k}, & \text{if } k \neq K \\ [\mathbf{w}_\phi^{K-k}, \mathbf{w}_\psi^{K-k}], & \text{if } k = K, \end{cases} \quad (6)$$

in which, $[\cdot, \cdot]$ is concatenation operation. It should be noted that the size for \mathbf{w}_ϕ^j is down-sampled as j increases during the fast DWT process. Besides, k is negative related to j , which means features from shallow layers (smaller k) correspond to coefficients of higher frequency (larger j). In this way, the target for MIM becomes reconstructing the wavelet coefficients of frequency $j = K - k$ with the features from layer l_k of the encoder.

Implementation Details

Our approach, called Wavelet-Driven Masked Image Modeling (WaMIM), is a simple yet powerful method for self-supervised pre-training. We leverage the wavelet transform tool to create multiple reconstruction targets representing different frequencies and scales. And multi-level features from different layers are utilized to predict these targets separately, ensuring that the model comprehensively learns rich image representations. Figure 2 shows the overview of our framework. We then sequentially detail the key components of our framework, including the encoder, decoder, masking strategy, reconstruction target, and training objectives.

Encoder. We adopt vision transformer and swin transformer as the foundational architectures for the encoder f_{EN} . The input image \mathbf{x} is first divided into regular non-overlapping patches, then some patches are masked out from the input, yielding a masked image $\hat{\mathbf{x}}$. For the vision transformer, only visible patches are fed into the encoder for efficiency. Thanks to the optimizations introduced by GreenMIM (Huang et al. 2022), we apply the same mechanism to the swin transformer as well. The input patches are embedded by a linear projection layer, incorporating positional embeddings, and then processed via a series of self-attention-based Transformer blocks. We sequentially select K features $\{z^{l_k} | k \in \{1, 2, \dots, K\}\}$, which comes from shallow to deep layers as the multi-level features outputted by the encoder. K equals the number of reconstruction targets. The encoder is further used for the various downstream tasks.

Decoder. We employ multiple decoders to predict different reconstruction targets using the encoded multi-level features. The decoders take the output features of the encoder along with mask tokens as input, where the mask token represents a shared and learnable vector. To mitigate the computational overhead resulting from the increase in decoder numbers and improve training efficiency, we adopt lightweight decoders, consisting of only one transformer block and a linear layer, for example. We denote the decoder attached to the feature of the l -th layer as f_{DE}^l .

Masking Strategy. We adopt the random block-wise masking strategy in our framework. Given that our method incor-

porates multiple reconstruction targets across various scales, it’s imperative to rescale the mask accordingly to match the corresponding scale of the reconstruction target during the computation of the reconstruction loss.

Reconstruction Targets. As shown in Section , we first perform multi-level wavelet decomposition on the input image to obtain multi-level wavelet coefficients. These wavelet coefficients are further utilized as multi-level reconstruction targets representing different frequencies and scales. For transformers, shallow layers generally learn low-level representations of images, *e.g.* textures and edges, while deeper layers capture higher-level semantic information, *e.g.* object shapes. Therefore, in our framework, the shallow-layer output features of the encoder are employed to predict high-frequency (low-level) reconstruction targets, whereas deep or final-layer output features are used to predict low-frequency (high-level) reconstruction targets. We denote the reconstruction target for l -th layer as $f_t^l(\mathbf{x})$.

Training objectives. The final reconstruction loss is defined by replacing the target generator in Equation 1 with wavelet coefficients at different scales, as expressed in

$$\mathcal{L} = \sum_{k=1}^{K-1} c^{l_k} \mathcal{D}(f_{AE}^{l_k}(f_m(\mathbf{x})), \mathbf{w}_\phi^{K-k}) + c^{l_K} \cdot \mathcal{D}(f_{AE}^{l_K}(f_m(\mathbf{x})), [\mathbf{w}_\psi^0, \mathbf{w}_\phi^0]), \quad (7)$$

where outputs of $f_{AE}^{l_k}$ will be made adaptive to different sizes of the target. Besides, the distance measurement $\mathcal{D}(\cdot, \cdot)$ encompasses masking operation considering different scales of the wavelet coefficients.

Experiments

Experimental settings

Pre-training setup. We perform pre-training on the ImageNet-1K dataset (Russakovsky et al. 2015) without any ground-truth labels. We use the columnar ViT (Dosovitskiy et al. 2021) and pyramidal Swin (Liu et al. 2021) architectures for the encoder, with an input size of 224×224 . The input images are segmented into patches of size $p = 16$ for ViT and $p = 4$ for Swin, and are randomly masked with a default ratio of $r = 0.75$. Basic data augmentations, including random cropping and horizontal flipping, are applied. For each architecture, we construct four reconstruction targets that vary from high-frequency, low-level to low-frequency, high-level. We perform a 5-level wavelet decomposition on the input image and select wavelet coefficients from levels 2 to 5 as the reconstruction targets, resulting in scales of $\{56^2, 28^2, 14^2, 7^2\}$. In the Swin architecture, we use output features from stages $\{2, 4, 22, 24\}$ for prediction. Each feature’s decoder consists of a transformer block with an embedding dimension of 128 and 4 attention heads. For ViT, the chosen layers are $\{3, 6, 9, 12\}$, with each decoder comprising a transformer block with an embedding dimension of 256 and 8 attention heads. Loss weights are set to $\{0.8, 0.9, 1.1, 1.2\}$. Haar wavelet basis is employed for wavelet transform, and the wavelet coefficients used for reconstruction are normalized. Detailed pre-training procedures are provided in the appendix.

Image Classification. After pre-training, we first conduct supervised end-to-end fine-tuning on ImageNet-1K (Russakovsky et al. 2015) dataset. The fine-tuning of all models is conducted using images with resolutions of 224×224 . The full fine-tuning details can be found in the appendix.

Object Detection and Instance Segmentation. We adopt the pre-trained Swin models as the feature extractor for Mask R-CNN (He et al. 2017). We start by fine-tuning Mask R-CNN with the COCO (Lin et al. 2014) train2017 dataset split, followed by an evaluation of its performance on the val2017 split using AP^b and AP^m metrics. The fine-tuning process is carried out using a $3\times$ schedule, consisting of 36 training epochs. We follow the implementation of Mask R-CNN provided by MMDetection (Chen et al. 2019). The full fine-tuning details are shown in the appendix.

Semantic Segmentation. We adapt the pre-trained ViT models for semantic segmentation using UperNet (Xiao et al. 2018) as the segmentor. We perform end-to-end fine-tuning on the ADE20k dataset (Zhou et al. 2017) for 160k iterations with an input resolution of 512×512 , and evaluate performance on the validation set using the mIoU metric. Full fine-tuning details are provided in the appendix.

Main Results

Image Classification. We evaluate our WaMIM against existing MIM models, analyzing both pre-training efficiency and top-1 fine-tuning accuracy. The results are illustrated in Figure 1c and summarized in Table 1. To ensure a fair comparison, we compute the pre-training efficiency of each method on identical hardware, utilizing a single Tesla A100-40G GPU, CUDA 11.7, and PyTorch 1.13.

For ViT-B, our method achieves 83.8% accuracy with only 400 epochs, 96 total GPU hours, which is 0.2% higher than the MAE at 1600 epochs while requiring only 13% of the computational cost. Regarding MFM, another method that employs the Fourier transform as the frequency analysis tool, our approach achieves a 0.7% higher fine-tuning accuracy while incurring only a 29% computational cost. This further highlights the superiority of wavelet transform over Fourier transform in MIM methods. When compared to LocalMIM, another multi-scale MIM method, our approach can also achieve better performance at both 100 epochs and 400 epochs, with higher training efficiency.

For the small-scale model ViT-S, our method can significantly enhance the performance even further. Our method at 100 epochs achieves a fine-tuning accuracy of 81.8%, which is 0.9%, 0.2%, and 0.2% higher than MAE, MFM, and LocalMIM at 300 epochs, despite using only 16%, 8%, and 25% computational costs. When enlarging the training schedule to 300 epochs, our method achieves an accuracy of 82.0%, marking a new state-of-the-art in terms of balanced performance and efficiency, to the best of our knowledge.

For the hierarchical Swin-B architecture, our method achieves a fine-tuning accuracy of 83.9% with only 100 epochs of pre-training, significantly outperforming SimMIM and GreenMIM in efficiency. Compared to LocalMIM, our method yields a slightly higher fine-tuning accuracy of 0.1% at 100 epochs, while requiring just 80% of the computational cost. With an extended training schedule of 400

Method	Model	Target	PT Epoch	GPU Hrs./Ep.	Total GPU Hrs.	Acc.
MAE* (He et al. 2022)	ViT-S	Pixel	300	0.32	96	80.9
LocalMIM* (Wang et al. 2023)	ViT-S	Hog	300	0.20	60	81.6
MFM (Xie et al. 2022a)	ViT-S	Fourier	300	0.60	180	81.6
WaMIM	ViT-S	Wavelet	100	0.15	15	81.8
WaMIM	ViT-S	Wavelet	300	0.15	45	82.0
MoCo v3 (Chen, Xie, and He 2021)	ViT-B	Momentum	600	-	-	83.2
iBOT (Zhou et al. 2022)	ViT-B	Momentum	400	4.33	1732	83.8
BEiT (Bao et al. 2021)	ViT-B	DALLE	800	1.03	824	83.2
MAE* (He et al. 2022)	ViT-B	Pixel	800	0.47	376	83.3
MAE (He et al. 2022)	ViT-B	Pixel	1600	0.47	752	83.6
MaskFeat(Wei et al. 2022)	ViT-B	Hog	1600	1.67	2672	84.0
CAE (Chen et al. 2024)	ViT-B	DALLE	800	1.20	960	83.6
LocalMIM (Wang et al. 2023)	ViT-B	Hog	100	0.30	30	83.3
LocalMIM* (Wang et al. 2023)	ViT-B	Hog	400	0.30	120	83.5
MFM (Xie et al. 2022a)	ViT-B	Fourier	300	1.10	330	83.1
WaMIM	ViT-B	Wavelet	100	0.24	24	83.4
WaMIM	ViT-B	Wavelet	400	0.24	96	83.8
SimMIM ₁₉₂ (Xie et al. 2022b)	Swin-B	Pixel	800	0.82	656	84.0
GreenMIM (Huang et al. 2022)	Swin-B	Pixel	800	0.37	296	83.7
LocalMIM (Wang et al. 2023)	Swin-B	Hog	100	0.50	50	83.8
LocalMIM (Wang et al. 2023)	Swin-B	Hog	400	0.50	200	84.1
WaMIM	Swin-B	Wavelet	100	0.40	40	83.9
WaMIM	Swin-B	Wavelet	400	0.40	160	84.1

Table 1: Top-1 accuracy (%) on ImageNet-1K. All models are pre-trained and fine-tuned at a resolution of 224×224 , except for SimMIM₁₉₂, which uses a 192×192 resolution during pre-training. * indicates results reproduced using the official code.

Method	Model	PT Ep.	PT Hrs.	AP ^b	AP ^m
SimMIM ₁₉₂	Swin-B	800	656	50.4	44.4
GreenMIM	Swin-B	800	296	50.0	44.1
LocalMIM	Swin-B	400	200	50.7	44.9
WaMIM	Swin-B	400	160	50.9	45.1

Table 2: The COCO results for object detection and instance segmentation with metrics including AP^b (%) and AP^m (%). All models are fine-tuned end-to-end using the Mask R-CNN framework with the pre-trained Swin-B backbone.

epochs, our method reaches an accuracy of 84.1%, solidifying its position as the top performer in both performance and efficiency compared to other methods.

Object Detection and Instance Segmentation. We evaluate our method’s transferability to downstream dense tasks, testing its performance on object detection and instance segmentation with the COCO dataset. Table 2 shows AP^b (object detection), AP^m (instance segmentation), and total pre-training hours across methods. Our method achieves gains of 0.5, 0.9, and 0.2 in object detection, and 0.7, 1.0, and 0.2 in instance segmentation, compared to SimMIM, GreenMIM, and LocalMIM, respectively, while incurring only 24%, 54%, and 80% of their computational costs.

Semantic Segmentation. We further evaluate our method on semantic segmentation tasks using the ADE20K dataset, reporting mIoU scores and total pre-training hours for different methods in Table 3. The results demonstrate that our method achieves comparable or superior performance while being significantly more efficient.

Method	Model	PT Ep.	PT Hrs.	mIoU
MoCo v3	ViT-B	-	-	47.3
BEiT	ViT-B	800	824	47.1
MAE	ViT-B	1600	752	48.1
MaskFeat	ViT-B	1600	2672	48.8
CAE	ViT-B	800	960	48.8
LocalMIM	ViT-B	400	120	47.2
WaMIM	ViT-B	400	96	48.7

Table 3: The ADE20K results for semantic segmentation in terms of mIoU. All models are fine-tuned end-to-end using UperNet with the pre-trained ViT-B backbone.

Ablation studies

In this part, we conduct ablation experiments to assess the impact of key components in our approach and validate the design choices. Table 4a~4f show the WaMIM ablation experimental results with ViT-B and Swin-B on ImageNet-1K. All models are pre-trained for 100 epochs. We report fine-tuning Top-1 accuracy (%). If not specified, the default setting is: the mask ratio is 75%, the decoder consists of a single Transformer block with a dimension of 256 and 8 heads for ViT and a dimension of 128 and 4 heads for Swin, the locations for output features are the output of each stage for Swin and the output of {3, 6, 9, 12}-th layer for ViT, the wavelet decomposition level is 5, the loss weight is {0.8, 0.9, 1.1, 1.2}. The selected settings are underlined.

Mask ratio. We experiment with mask ratios ranging from 0.4 to 0.9 in our method, as shown in Table 4a. The results indicate that increasing the mask ratio from 0.4 to 0.75 im-

ratio	ViT-B	Swin-B
0.4	83.1	83.6
0.75	83.4	83.9
0.9	82.9	83.4

(a)

levels	ViT-B	Swin-B
5	83.4	83.9
4	83.4	83.9
3	83.2	83.7

(b)

decoder	ViT-B	Swin-B
512D-16H	83.4	83.8
256D-8H	83.4	83.9
128D-4H	83.2	83.9

(c)

locations	model	Top-1
{12}		82.9
{1, 2, 11, 12}	ViT-B	83.2
{2, 4, 10, 12}		83.3
{3, 6, 9, 12}		83.4

(d)

locations	model	Top-1
{24}		83.2
{22, 24}	Swin-B	83.5
{4, 22, 24}		83.8
{2, 4, 22, 24}		83.9

(e)

loss weight	ViT-B	Swin-B
{1, 1, 1, 1}	83.3	83.8
{0.8, 0.9, 1.1, 1.2}	83.4	83.9
{0.7, 0.8, 1.2, 1.3}	83.2	83.7
{1.2, 1.1, 0.9, 0.8}	83.1	83.5

(f)

Table 4: WaMIM ablations. (a) **Mask ratio**. A mask ratio of 0.75 performs best for both ViT-B and Swin-B. (b) **Wavelet decomposition levels**. Five-level wavelet decomposition demonstrates the best performance and efficiency. (c) **Decoder design**. The light-weight decoder matches the heavier one in performance with greater efficiency. (d) **Locations for ViT**. The uniform division for chosen layers achieves the best performance for ViT. (e) Predicting multi-frequency, multi-scale targets boosts performance and efficiency. (f) **Loss weight**. Emphasizing low-frequency, high-level targets improves representation learning.

proves performance on ViT-B and Swin-B. However, at a high mask ratio of 0.9, performance declines, suggesting that the reconstruction task becomes too difficult for the model to learn. Thus, we selected 0.75 as the default mask ratio.

Decoder design. Numerous studies (Huang et al. 2022; Li et al. 2022) show that light-weight decoders effectively enable MIM to acquire generalizable representations. In our framework, we create multiple reconstruction targets at different frequencies and scales, requiring several independent decoders. To reduce computational costs, we explore the use of lighter decoders. This section examines how the embedding dimension and the number of self-attention heads in a single-transformer-block decoder affect performance. As shown in Table 4c, our approach achieves comparable or better results, even with significantly lighter decoders.

Wavelet decomposition levels. Wavelet decomposition levels determine the frequency and scale granularity represented by each level of wavelet coefficients. For example, a 5-level decomposition on an image yields coefficients across 6 frequency bands with corresponding scales of $\{112^2, 56^2, 28^2, 14^2, 7^2\}$. We experiment with 3, 4, and 5 decomposition levels, setting the number of reconstruction targets to 3, 4, and 4, respectively. We then select coefficients from the $\{1, 2, 3\}$ -th, $\{1, 2, 3, 4\}$ -th, and $\{2, 3, 4, 5\}$ -th levels, corresponding to scales of $\{112^2, 56^2, 28^2\}$, $\{112^2, 56^2, 28^2, 14^2\}$, and $\{56^2, 28^2, 14^2, 7^2\}$, ranging from high to low frequency. As shown in Table 4b, the 5-level wavelet decomposition offers the best balance of efficiency and performance.

Loss weight. Loss weights determine whether we prioritize reconstructing high-frequency, low-level targets or low-frequency, high-level targets during representation learning. Table 4f shows our method’s performance under different loss weight configurations. Lower indices in the array correspond to weights for high-frequency, low-level targets, while higher indices represent weights for low-frequency, high-level targets. The results suggest that moderately prioritizing low-frequency, high-level features enhances representation learning. However, overemphasizing them can hinder

performance, underscoring the importance of incorporating high-frequency, low-level information in MIM.

Locations for multi-level features. In our framework, the locations of multi-level features used to predict reconstruction targets is crucial because each transformer layer focuses on different image frequencies and scales. Shallow layers primarily capture high-frequency, low-level details, while deep layers capture low-frequency, high-level semantic information. Thus, selecting the appropriate features to match the frequency and scale of reconstruction targets can significantly improve representation learning efficiency and performance. For the hierarchical Swin architecture, we select output features from its four stages, corresponding to the $\{2, 4, 22, 24\}$ -th layers from shallow to deep. In our experiments, we progressively add lower-stage features, starting from the last stage, to demonstrate the benefit of predicting multiple reconstruction targets across different frequencies and scales. As shown in Table 4e, predicting multiple carefully designed targets enhances both efficiency and performance. For the columnar ViT, we test various layer combinations, focusing on $\{1, 2, 11, 12\}$ -th (head and tail division), $\{2, 4, 10, 12\}$ -th (Swin-style division), and $\{3, 6, 9, 12\}$ -th (uniform division). The results in Table 4d indicate that uniform division yields the best performance for ViT.

Conclusion

In this paper, we introduce an efficient approach to enhance masked image modeling (MIM) by leveraging wavelet transform to guide the construction of multi-level reconstruction targets representing different frequencies and scales. We employ multi-level features to predict these reconstruction targets separately. In practice, we directly utilize multi-level wavelet coefficients as reconstruction targets, ensuring that our method is both straightforward and efficient. Moreover, our approach can seamlessly integrate into most existing MIM methods. Extensive experimental results indicate that, compared to other methods, our approach achieves comparable or superior results on numerous downstream tasks, while also demonstrating significantly higher efficiency.

Acknowledgments

This work was supported in part by the National Natural Science Foundation under Grant 62402251, 62472238.

References

- Arts, L. P.; and van den Broek, E. L. 2022. The fast continuous wavelet transformation (fCWT) for real-time, high-quality, noise-resistant time–frequency analysis. *Nature Computational Science*, 2(1): 47–58.
- Bao, H.; Dong, L.; Piao, S.; and Wei, F. 2021. BEiT: BERT Pre-Training of Image Transformers. In *International Conference on Learning Representations*.
- Chen, K.; Wang, J.; Pang, J.; Cao, Y.; Xiong, Y.; Li, X.; Sun, S.; Feng, W.; Liu, Z.; Xu, J.; et al. 2019. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*.
- Chen, X.; Ding, M.; Wang, X.; Xin, Y.; Mo, S.; Wang, Y.; Han, S.; Luo, P.; Zeng, G.; and Wang, J. 2024. Context autoencoder for self-supervised representation learning. *International Journal of Computer Vision*, 132(1): 208–223.
- 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*, 9640–9649.
- Dalal, N.; and Triggs, B. 2005. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, volume 1, 886–893. Ieee.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Fang, Y.; Dong, L.; Bao, H.; Wang, X.; and Wei, F. 2022. Corrupted image modeling for self-supervised visual pre-training. *arXiv preprint arXiv:2202.03382*.
- Fuoli, D.; Van Gool, L.; and Timofte, R. 2021. Fourier space losses for efficient perceptual image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2360–2369.
- Gao, P.; Ma, T.; Li, H.; Lin, Z.; Dai, J.; and Qiao, Y. 2022. Convmae: Masked convolution meets masked autoencoders. *arXiv preprint arXiv:2205.03892*.
- 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.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- Huang, L.; You, S.; Zheng, M.; Wang, F.; Qian, C.; and Yamazaki, T. 2022. Green hierarchical vision transformer for masked image modeling. *Advances in Neural Information Processing Systems*, 35: 19997–20010.
- Jiang, L.; Dai, B.; Wu, W.; and Loy, C. C. 2021. Focal frequency loss for image reconstruction and synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13919–13929.
- Kenton, J. D. M.-W. C.; and Toutanova, L. K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*, 4171–4186.
- Li, X.; Wang, W.; Yang, L.; and Yang, J. 2022. Uniform masking: Enabling mae pre-training for pyramid-based vision transformers with locality. *arXiv preprint arXiv:2205.10063*.
- Li, X.; Zhang, Y.; Yuan, J.; Lu, H.; and Zhu, Y. 2023. Discrete Cosin TransFormer: Image Modeling From Frequency Domain. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5468–5478.
- 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, C.; Xiang, W.; He, Y.; Xue, H.; Zheng, S.; and Su, H. 2023a. Improving model generalization by on-manifold adversarial augmentation in the frequency domain. *arXiv preprint arXiv:2302.14302*.
- Liu, H.; Jiang, X.; Li, X.; Guo, A.; Hu, Y.; Jiang, D.; and Ren, B. 2023b. The devil is in the frequency: Geminated gestalt autoencoder for self-supervised visual pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1649–1656.
- Liu, T.; Cheng, J.; and Tan, S. 2023. Spectral Bayesian Uncertainty for Image Super-Resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18166–18175.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Liu, Y.; Zhang, S.; Chen, J.; Chen, K.; and Lin, D. 2023c. PixMIM: Rethinking Pixel Reconstruction in Masked Image Modeling. *arXiv preprint arXiv:2303.02416*.
- Liu, Y.; Zhang, S.; Chen, J.; Yu, Z.; Chen, K.; and Lin, D. 2023d. Improving pixel-based mim by reducing wasted modeling capability. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5361–5372.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Park, J.; and Johnson, J. 2023. RGB no more: Minimally-decoded JPEG Vision Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22334–22346.
- Park, N.; and Kim, S. 2022. How do vision transformers work? *arXiv preprint arXiv:2202.06709*.

- Phung, H.; Dao, Q.; and Tran, A. 2023. Wavelet diffusion models are fast and scalable image generators. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10199–10208.
- Ren, S.; Wei, F.; Albanie, S.; Zhang, Z.; and Hu, H. 2023. DeepMIM: Deep Supervision for Masked Image Modeling. *arXiv preprint arXiv:2303.08817*.
- 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*, 115: 211–252.
- Said, S.; Jemai, O.; Hassairi, S.; Ejbali, R.; Zaied, M.; and Ben Amar, C. 2016. Deep wavelet network for image classification. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 000922–000927.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Wang, H.; Tang, Y.; Wang, Y.; Guo, J.; Deng, Z.-H.; and Han, K. 2023. Masked Image Modeling with Local Multi-Scale Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2122–2131.
- 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.
- Xiang, W.; Liu, C.; Su, H.; and Yu, H. 2024. AEMIM: Adversarial Examples Meet Masked Image Modeling. *arXiv preprint arXiv:2407.11537*.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified perceptual parsing for scene understanding. In *Proceedings of the European conference on computer vision (ECCV)*, 418–434.
- Xie, J.; Li, W.; Zhan, X.; Liu, Z.; Ong, Y. S.; and Loy, C. C. 2022a. Masked frequency modeling for self-supervised visual pre-training. *arXiv preprint arXiv:2206.07706*.
- Xie, Z.; Zhang, Z.; Cao, Y.; Lin, Y.; Bao, J.; Yao, Z.; Dai, Q.; and Hu, H. 2022b. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9653–9663.
- Xu, K.; Qin, M.; Sun, F.; Wang, Y.; Chen, Y.-K.; and Ren, F. 2020. Learning in the frequency domain. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1740–1749.
- 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.
- Yao, T.; Pan, Y.; Li, Y.; Ngo, C.-W.; and Mei, T. 2022. Wavevit: Unifying wavelet and transformers for visual representation learning. In *European Conference on Computer Vision*, 328–345. Springer.
- You, Z.; Liu, D.; Han, B.; and Xu, C. 2024. Beyond pre-trained features: Noisy image modeling provides adversarial defense. *Advances in Neural Information Processing Systems*, 36.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 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*.
- Zhu, Z.; and Soricut, R. 2024. Wavelet-Based Image Tokenizer for Vision Transformers. *arXiv preprint arXiv:2405.18616*.