

MAGE: MAsked Generative Encoder to Unify Representation Learning and Image Synthesis

Tianhong Li^{1*} Huiwen Chang³ Shlok Kumar Mishra² Han Zhang³
Dina Katabi¹ Dilip Krishnan³
¹MIT CSAIL ²University of Maryland ³Google Research

Abstract

Generative modeling and representation learning are two key tasks in computer vision. However, these models are typically trained independently, which ignores the potential for each task to help the other, and leads to training and model maintenance overheads. In this work, we propose MAsked Generative Encoder (MAGE), the first framework to unify SOTA image generation and self-supervised representation learning. Our key insight is that using variable masking ratios in masked image modeling pre-training can allow generative training (very high masking ratio) and representation learning (lower masking ratio) under the same training framework. Inspired by previous generative models, MAGE uses semantic tokens learned by a vector-quantized GAN at inputs and outputs, combining this with masking. We can further improve the representation by adding a contrastive loss to the encoder output. We extensively evaluate the generation and representation learning capabilities of MAGE. On ImageNet-1K, a single MAGE ViT-L model obtains **9.10 FID** in the task of class-unconditional image generation and **78.9% top-1 accuracy** for linear probing, achieving state-of-the-art performance in both image generation and representation learning. Code is available at <https://github.com/LTH14/mage>.

1. Introduction

In recent years, we have seen rapid progress in both generative models and representation learning of visual data. Generative models have demonstrated increasingly spectacular performance in generating realistic images [3, 7, 15, 46], while state-of-the-art self-supervised representation learning methods can extract representations at a high semantic level to achieve excellent performance on a number of

*This work was done when Tianhong Li interned at Google Research, and was partially supported by the MIT-IBM Watson Research Collaboration grant. Correspondence to: Tianhong Li <tianhong@mit.edu>, Huiwen Chang <huiwenchang@google.com>.




Figure 1. Linear probing and class unconditional generation performance of different methods trained and evaluated on ImageNet-1K. MAGE achieves SOTA performance in linear probing and establishes a new SOTA in class unconditional generation.

downstream tasks such as linear probing and few-shot transfer [2, 6, 8, 13, 25, 26].

Currently, these two families of models are typically trained independently. Intuitively, since generation and recognition tasks require both visual and semantic understanding of data, they should be complementary when combined in a single framework. Generation benefits representation by ensuring that both high-level semantics and low-level visual details are captured; conversely, representation benefits generation by providing rich semantic guidance. Researchers in natural language processing have observed this synergy: frameworks such as BERT [14] have both high-quality text generation and feature extraction. Another example is DALLE-2 [43], where latents conditioned on a *pre-trained* CLIP representation are used to create high-quality text-to-image generations.

However, in computer vision, there are currently no widely adopted models that unify image generation and representation learning in the same framework. Such uni-



Figure 2. Reconstruction results using MAE and MAGE with 75% masking ratio. MAE reconstructs blurry images with low quality, while MAGE can reconstruct high-quality images with detail, and further improves quality through iterative decoding (see subsection 3.2 for details). With the same mask, MAGE generates diverse reconstruction results with different random seeds. Note that the mask for MAGE is on semantic tokens whereas that of MAE is on patches in the input image.

fication is non-trivial due to the structural difference between these tasks: in generative modeling, we **output** high-dimensional data, conditioned on low-dimension inputs such as class labels, text embeddings, or random noise. In representation learning, we **input** a high-dimensional image and create a low-dimensional compact embedding useful for downstream tasks.

Recently, a number of papers have shown that representation learning frameworks based on masked image modeling (MIM) can obtain high-quality representations [2, 18, 26, 31], often with very high masking ratios (e.g. 75%) [26]. Inspired by NLP, these methods mask some patches at the input, and the pre-training task is to reconstruct the original image by predicting these masked patches. After pre-training, task-specific heads can be added to the encoder to perform linear probe or fine-tuning.

These works inspire us to revisit the unification question. Our key insight in this work is that generation is viewed as “reconstructing” images that are 100% masked, while representation learning is viewed as “encoding” images that are 0% masked. We can therefore enable a unified architecture by using a *variable masking ratio* during MIM pre-training. The model is trained to reconstruct over a *wide range* of masking ratios covering high masking ratios that enable generation capabilities, and lower masking ratios that enable representation learning. This simple but very effective approach allows a smooth combination of generative training and representation learning in the *same framework*: same architecture, training scheme, and loss function.

However, directly combining existing MIM methods

with a variable masking ratio is insufficient for high quality generation because such methods typically use a simple reconstruction loss on pixels, leading to blurry outputs. For example, as a representative of such methods, the reconstruction quality of MAE [27] is poor: fine details and textures are missing (Figure 2). A similar issue exists in many other MIM methods [11, 36].

This paper focuses on bridging this gap. We propose MAGE, a framework that can both generate realistic images and extract high-quality representations from images. Besides using variable masking ratio during pre-training, unlike previous MIM methods whose inputs are pixels, both the inputs and the reconstruction targets of MAGE are *semantic tokens*. This design improves both generation and representation learning, overcoming the issue described above. For generation, as shown in Figure 2, operating in token space not only allows MAGE to perform image generation tasks iteratively (subsection 3.2), but also enables MAGE to learn a probability distribution of the masked tokens instead of an average of all possible masked pixels, leading to diverse generation results. For representation learning, using tokens as inputs and outputs allows the network to operate at a high semantic level without losing low-level details, leading to significantly higher linear probing performances than existing MIM methods.

We evaluate MAGE on multiple generative and representation downstream tasks. As shown in Figure 1, for class-**un**conditional image generation on ImageNet-1K, our method surpasses state of the art with both ViT-B and ViT-L (ViT-B achieves 11.11 FID [29] and ViT-L achieves 9.10

FID), outperforming the previous state-of-the-art result by a large margin (MaskGIT [7] with 20.68 FID). This significantly push the limit of class-unconditional generation to a level even close to the state-of-the-art of class-conditional image generation (~ 6 FID [7, 46]), which is regarded as a much easier task in the literature [38]. For linear probing on ImageNet-1K, our method with ViT-L achieves 78.9% top-1 accuracy, surpassing all previous MIM-based representation learning methods and many strong contrastive baselines such as MoCo-v3 [13]. Moreover, when combined with a simple contrastive loss [9], MAGE-C with ViT-L can further get 80.9% accuracy, achieving state-of-the-art performance in self-supervised representation learning. We summarize our contributions:

- We introduce MAGE, a novel method that unifies generative model and representation learning by a single token-based MIM framework with variable masking ratios, introducing new insights to resolve the unification problem.
- MAGE establishes a new state of the art on the task of class-unconditional image generation on ImageNet-1K.
- MAGE further achieves state of the art in different downstream tasks, such as linear probing, few-shot learning, transfer learning, and class-conditional image generation.

2. Related Work

Self-supervised Learning in Computer Vision. Early work on unsupervised representation learning has focused on designing pretext tasks and training the network to predict their pseudo labels. Such tasks include solving jigsaw puzzles [39], restoring a missing patch [41], or predicting image rotation [23]. These pretext tasks result in representations that significantly trailed supervised training.

Contrastive learning [8, 10, 34, 40] has proven to be a competitive and systematic method to learn effective representations without human supervision, getting performance very close to that of supervised pre-training. A number of variants of contrastive learning have been proposed: SimCLR [8] uses a large batch size, and samples negative pairs within each batch; momentum-contrastive approach (MoCo) [28] leverages a moving-average encoder and a queue to store negative samples during training; Contrastive-Multiview-Coding [50] maintains a memory-bank to store features and generate negative samples. Some recent methods, like BYOL, do not rely on negative pairs [12, 25]. Instead, they use two neural networks that learn from each other to boost performance.

Recently, vision researchers have found that masked image modeling (MIM), modeled after techniques in NLP e.g. [14], is a very effective task for self-supervised learning. BEiT [2] recovers discrete visual tokens from masked inputs. PeCo [18] further regards MoCo-v3 [13] as the perceptual model in VQGAN training to get a better tokenizer.

MAE [26] considers MIM as a denoising pixel-level reconstruction task, and CMAE [31] further combines MAE with a contrastive loss. Some other methods such as MaskFeat [52] and MVP [53] predict features generated from teacher models.

However, current self-supervised learning methods based on MIM favor the performance of the representations on downstream tasks instead of the quality of the reconstructed images, leading to poor reconstructive results [2, 26]. Our paper for the first time shows that a single model can not only learn high-level fine-grained representations, but also be used to generate images of high visual fidelity.

Generative Models for Image Synthesis. Recent years have witnessed tremendous progress in deep generative models for image synthesis. One major stream of generative models is built on top of generative adversarial networks (GANs) [4, 24, 32, 56, 57]. GAN-based models can generate realistic images in various domains, but suffer from training instabilities and mode collapse issues. Another stream is based on a two-stage scheme [7, 33, 44, 51, 55]: first tokenize the image into a latent space and then apply maximum likelihood estimation and sampling in the latent space. VQVAE-2 [44] first shows this two-stage scheme can generate more diverse samples than GANs. ViT-VQGAN [55] uses ViT-based [19] encoder and decoder to get the latent code and then apply autoregressive generation in the latent space. MaskGIT [7] explores using a bidirectional transformer for token modeling and proposes parallel decoding for faster inference speed. Very recently, diffusion models [15, 30, 46, 49] have also achieved superior results on image synthesis.

However, the above generative models lack the ability to extract high-quality semantic representations from images. Prior works [16, 17, 21, 54, 55] explore the possibility of using latent features as representations, but their performance is sub-optimal. Our method surpasses previous generative models on both class unconditional generation and representation learning by a large margin, showing that a unified, high-performance framework is feasible.

3. Method

MAGE is a unified framework for both generative tasks and representation learning. To enable such unification, we first use a pre-trained VQGAN model [22] to quantize input images into semantic tokens. Then we randomly mask out some input tokens using a variable masking ratio ranging from 0.5 to 1 (see Figure 3), and apply an encoder-decoder transformer architecture on the remaining (unmasked) tokens to predict the masked tokens. We can further improve the separability of the learned representation by adding a simple yet effective contrastive loss similar to SimCLR [9] on the output of the encoder (MAGE-C). Below, we de-

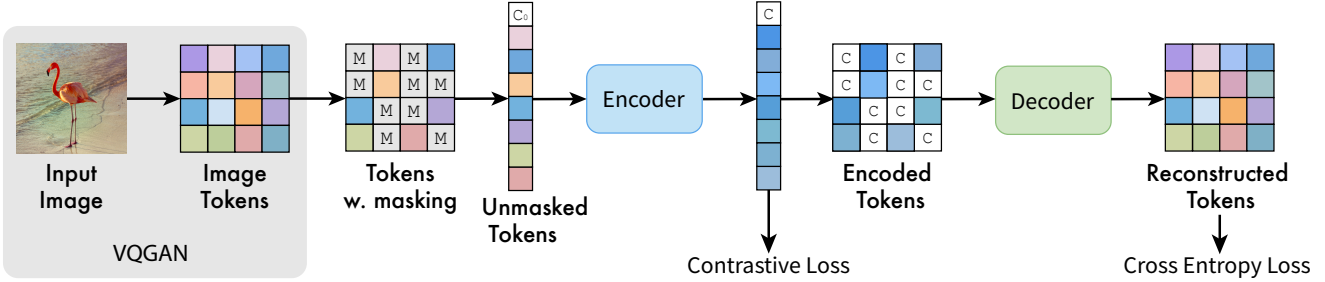


Figure 3. MAGE Framework: we first use a VQGAN tokenizer to tokenize the input image into a sequence of semantic tokens. We then sample a masking ratio (see text for details on the sampling strategy) and randomly mask out tokens according to this sampled ratio. A ViT encoder-decoder structure processes the unmasked tokens. A reconstructive cross-entropy loss encourages the model to reconstruct masked tokens. We can also add an optional contrastive loss at the output of the encoder to further improve the linear separability of the learned latent feature space.

scribe our design in detail.

3.1. Pre-training

Tokenization. We first tokenize the input image into a sequence of semantic tokens using a tokenizer. The tokenizer employs the same setup as the first stage in the VQGAN model [22]. This tokenization step allows our model to operate on semantic tokens instead of raw pixels, which is beneficial for both generation and representation learning as shown in Figure 2 and Figure 6.

Masking Strategy. To further bridge the gap between generative modeling and representation learning, we adopt a masking strategy with variable masking ratios. Specifically, we first randomly sample the masking ratio m_r from a truncated Gaussian distribution centered at 0.55, left truncated by 0.5, and right truncated by 1. If the length of the input sequence of tokens is l , we randomly mask out $m_r \cdot l$ tokens and replace them with a learnable mask token [M] (Figure 3). Since $m_r \geq 0.5$, we further randomly drop out $0.5 \cdot l$ tokens from those masked tokens. Dropping a large fraction of masked tokens significantly reduces overall pre-training time and memory consumption, while helping both generation and representation performance. This is consistent with the findings in MAE [26] for representation performance.

Encoder-Decoder Design. After masking and dropping input tokens, following [20], we concatenate a learnable “fake” class token [C₀] to the input sequence. The concatenated sequence is then fed into a Vision Transformer (ViT) [20] encoder-decoder structure. Specifically, the ViT encoder takes the sequence of tokens after masking and dropping as input and encodes them into latent feature space. Before decoding, the output of the encoder is first padded to the full input length using the class token feature [C] learned by the encoder. As shown in MAE [26], the class token position can summarize global features of the input image. Thus, instead of using a learnable masking token that is shared across different images, we use [C] that is specific to each image to pad the encoder outputs. We

show in the Appendix that this design improves both generation and representation learning performance over using a masking token (as done in MAE [26]). The decoder then takes the padded features to reconstruct the original tokens.

Reconstructive Training. Let $Y = [y_i]_{i=1}^N$ denote the latent tokens obtained from the tokenizer, where N is the token sequence length, and $M = [m_i]_{i=1}^N$ denotes a corresponding binary mask determining which tokens are to be masked. The training objective is to reconstruct the masked tokens from the unmasked tokens. Therefore, we add a cross-entropy loss between the ground-truth one-hot tokens and the output of the decoder. Specifically,

$$\mathcal{L}_{reconstructive} = -\mathbb{E}_{Y \in \mathcal{D}} \left(\sum_{\forall i, m_i=1} \log p(y_i | Y_M) \right), \quad (1)$$

where Y_M are the (subset of) *unmasked* tokens in Y and $p(y_i | Y_M)$ is the probability predicted by the encoder-decoder network, conditioned on the unmasked tokens. Following MAE, we only optimize this loss on *masked* tokens (optimizing the loss on all tokens reduces both generation and representation learning performance, similar to the observations in [26]).

Contrastive Co-training. As shown in [35] and [31], adding a contrastive loss in MIM method can further improve its representation learning performance. In our MAGE framework, we can also add a contrastive loss to force better linear separability of the learned feature space. Similar to SimCLR [10], we add a two-layer MLP on top of the feature obtained by globally average pooling the output of the encoder. We then add an InfoNCE loss [40] on the output of the MLP head:

$$\mathcal{L}_{contrastive} = -\frac{1}{B} \sum_{i=1}^B \log \frac{e^{z_i^T \cdot z_i^+ / \tau}}{\sum_{j=1}^B e^{z_i^T \cdot z_j / \tau}}, \quad (2)$$

where z denotes the normalized features after the two-layer MLP, B denotes the batch size, and τ denotes the temperature. The positive pairs z_i, z_i^+ are from two augmented

views of the same image, and the negative samples z_j are all other samples in the same batch. Our final loss is:

$$\mathcal{L} = \mathcal{L}_{reconstructive} + \lambda \cdot \mathcal{L}_{contrastive} \quad (3)$$

where $\lambda = 0.1$ balances the scale of the two losses. We do not use the extensive augmentations typically used in contrastive learning, such as color jitter, random grey scale, or gaussian noise. This is because the reconstructive loss acts as a regularizer that prevents the encoder from learning shortcut solutions [45]. Our approach achieves superior performance on both generative tasks and representation learning even without the contrastive loss, and representation learning performance can be further boosted with the contrastive loss.

3.2. Post-training Evaluation

To generate images for generative model evaluation, we use a **iterative decoding** strategy similar to MaskGIT [7]. We start from a blank image with all the tokens masked out. At each iteration, our model first predicts the tokens for the remaining masked tokens. Then we sample some of the predicted tokens (tokens that have a higher predicted probability are of higher probability to be sampled) and replace the corresponding masked tokens with these sampled predicted tokens. The number of masked tokens to be replaced in each iteration follows a cosine function, i.e., we replace fewer masked tokens in the early iterations and more masked tokens at later iterations. We use a total of 20 steps to generate an image. For representation learning, we globally average pool the features output from the ViT encoder, and use the pooled features as the input features for the classification head. A detailed description of our pre-training and evaluation implementations and architectures is provided in the Appendix.

4. Results

MAGE is a unified framework for both generative model and representation learning. In this section, we conduct extensive experiments to evaluate the generation as well as visual representation capabilities. To evaluate MAGE’s generative performance, we conduct experiments on ImageNet-1K dataset [47] for the task of class-unconditional image generation. To evaluate the quality of the learned representations, we conduct experiments on ImageNet-1K dataset [47] under two protocols: first is linear probing, where we add a linear classification head on top of the learned representations and only train the classification head, while keeping the backbone frozen; second is fine-tuning, where we fine-tune the whole parameters for the classification task. We also include results on few-shot learning and transfer learning to better evaluate the quality of the representations. More results and ablation studies can be found in the Appendix.

4.1. Pre-training Setup

We set the input image resolution as 256x256 to be consistent with previous generative models. After passing through the VQGAN tokenizer, the token sequence length is 16x16 (256 tokens). Following MAE [26], we use strong random crop and resize (0.2 to 1) and random flipping as our default augmentations. We also trained models with a weaker version of random crop and resize (range from 0.8 to 1), which we call “w.a.” in the results. We pre-train base- and large-size vision Transformers [20], i.e., ViT-B and ViT-L, respectively. We use AdamW to train the model for 1600 epochs with batch size of 4096 for ViT-B, and batch size of 2048 for ViT-L. We use a cosine learning rate schedule with an 80-epoch warmup. The base learning rate is 1.5×10^{-4} for both ViT-B and ViT-L, and is further scaled by $batchsize/256$. More details are in the Appendix.

4.2. Image Generation

Table 1. Quantitative comparison with state-of-the-art generative models on ImageNet 256x256 for class-unconditional generation. We train LDM-8 [46] on ImageNet by ourselves using the official codebase. All baselines are reported without classifier guidance as it requires class label during training. Classifier-free guidance cannot improve unconditional generation because the guidance itself is generated by unconditional generation.

Methods	RES	FID↓	IS↑	#params
Self-Conditioned GAN [37]	128	40.3	15.82	-
BigGAN [17]	256	38.6	24.70	~70M
BigGAN [17]	128	30.9	23.56	~70M
BigGAN+Clustering [38]	128	22.0	23.5	~70M
HiT [58]	128	30.8	21.64	~30M
LDM [46]	256	39.1	22.83	395M
ADM [15]	256	26.2	39.70	554M
MaskGIT [7]	256	20.7	42.08	227M
MAGE (ViT-B)	256	11.1	81.17	200M
MAGE (ViT-B, w.a.)	256	8.67	94.8	200M
MAGE (ViT-L)	256	9.10	105.1	463M
MAGE (ViT-L, w.a.)	256	7.04	123.5	463M

Class-Unconditional Image Generation. Our pre-trained model can naturally perform class-unconditional image generation without any fine-tuning on the model parameters. Table 1 compares the class-unconditional image generation results of our model and SOTA generative models on ImageNet, reporting Fréchet Inception Distance (FID) [29] and Inception Score (IS) [48] as standard metrics. As shown in the table, our method outperforms all previous image generation methods by a large margin. The previous SOTA can only achieve 20.7 FID and 42.08 IS, while our ViT-B model can achieve 11.1 FID and 81.17 IS with similar number of parameters. This is likely because our framework can extract much better representations than all previous generative models as shown in Ta-



(a) Default augmentation (b) Weak augmentation

Figure 4. Images generated by MAGE (ViT-L). (a) images generated from MAGE trained with default strong augmentation, i.e., crops out larger portion of the image. (b) images generated from MAGE trained with weak augmentations, i.e., crops out smaller portion of the image. We see that visual fidelity and diversity are very good for both models.

ble 2, leading to superior generative performance. Our ViT-L model further achieves 9.10 FID and 7.04 FID when trained with weak augmentation, which is very close to the *class-conditional* generation performance of transformer models (e.g. 6.18 FID in MaskGIT [7]), a much easier task than class-unconditional generation [38].

We also note that the augmentations used to train the tokenizer and the encoder-decoder model can affect the evaluation scores. As shown in Table 1 and Figure 4, with default “strong” augmentation (i.e. random resized crop scale from 0.2 to 1), the FID and IS of the model are worse than using a weaker augmentation (random resized crop scale from 0.8 to 1). One possible reason is that the ImageNet validation set used to compute the FID is resized to 256 and center cropped. Since FID is computed based on the similarity between the generated image and the images in ImageNet validation set, it will be higher if the scale of the generated image is smaller. However, this does not necessarily mean that the visual quality of the generated image is worse. As shown in Figure 4, images generated with default augmentation can be much more zoomed in and off-center, but the images are still realistic and of high quality. We include more results in the Appendix, including for class conditional generation and image editing tasks such as inpainting.

4.3. Image Classification

Linear Probing. Linear probing is a primary evaluation protocol for self-supervised learning. As shown in Table 2, MAGE outperforms MAE [26] by 6.7% on ViT-B and 3.1% on ViT-L for ImageNet-1K linear probe top-1 accuracy, achieving state-of-the-art results among all MIM methods. Moreover, a simple contrastive loss similar to SimCLR [8] can further boost our performance. We do not use color jitter, random grey scale, or multi-crop augmentations used in SwAV [5], DINO [5] and iBOT [59]. Multi-crop augmentation typically brings 3%-5% improvements on accuracy,

Table 2. Top-1 accuracy of linear probing on ImageNet-1k. † denotes methods which require additional teacher model (CLIP) trained from image-text data. * denotes methods using multi-crop augmentations. RN is short for ResNet. The number of parameters for MAGE includes VQ-GAN tokenizer and ViT encoder.

Methods	Model	#params	Acc.
<i>Generative models</i>			
BigBiGAN [17]	RN50	23M	56.6
MaskGIT [7]	BERT	227M	57.4
ViT-VQGAN [55]	VIM-Base	650M	65.1
ViT-VQGAN [55]	VIM-Large	1697M	73.2
<i>MIM methods</i>			
BEiT [2]	ViT-B	86M	56.7
MAE [26]	ViT-B	86M	68.0
Ge ² -AE [36]	ViT-B	86M	75.3
MAGE	ViT-B	24M+86M	74.7
MAE [26]	ViT-L	304M	75.8
MAGE	ViT-L	24M+304M	78.9
<i>Contrastive methods</i>			
SimCLRv2 [10]	RN50w2	94M	75.6
BYOL [25]	RN50w2	94M	77.4
CAE [11]	ViT-B	86M	70.4
CMAE [31]	ViT-B	86M	73.9
MoCo v3 [13]	ViT-B	86M	76.7
DINO [59]	ViT-B	86M	72.8
iBOT [59]	ViT-B	86M	76.0
MAGE-C	ViT-B	24M+86M	78.2
SimCLRv2 [10]	RN152w2	233M	77.4
BYOL [25]	RN200w2	250M	79.6
MoCo v3 [13]	ViT-L	304M	77.6
CAE [11]	ViT-L	304M	78.1
MoCo v3 [13]	ViT-H	632M	78.1
MAGE-C	ViT-L	24M+304M	80.9
<i>Additional Data/Aug.</i>			
MVP [†] [53]	ViT-B	86M	75.4
BEiT v2 [†] [42]	ViT-B	86M	80.1
SwAV* [5]	RN50w5	586M	78.5
DINO* [6]	ViT-B	86M	78.2
iBOT* [59]	ViT-B	86M	79.5
iBOT* [59]	ViT-L	304M	81.0

but introduces large computational overheads. In spite of no multi-crop, MAGE-C achieves 78.2% accuracy with ViT-B and 80.9% accuracy with ViT-L. Our ViT-B performance surpasses that of ViT-H in MoCo v3 (632M parameters), indicating that the extra parameters (24M) in the tokenizer are *not* the reason for our good performance.

Few-shot Learning. The premise of self-supervised learning is to learn representations on unlabeled data that can be effectively applied to prediction tasks with few labels [10]. Following [19], we freeze the weights of the pre-trained model and train a linear classifier on top using a few labeled samples. As shown in Table 3, our methods with ViT-B outperform MAE [26] by a very large margin




Figure 5. Transfer learning performance of ViT-B and ViT-L pre-trained on ImageNet-1K using different methods. Our methods outperforms SimCLR [9] and MAE [26] on 6 of the 8 datasets.

Table 3. Few-shot evaluation on ImageNet-1K. We report the top-1 accuracy with different self-supervised methods and different numbers of the ImageNet-1K labels used. We report the accuracy of MAE under our implementation (denoted by †). Note that MSN [1] uses multi-crop augmentation.

Method	Training images per ImageNet Class			
	5	10	13	25
<i>ViT-B</i>				
MAE† [26]	29.2	34.5	-	38.7
MSN [1]	65.5	-	69.6	-
MAGE	53.5	58.4	59.7	61.7
MAGE-C	62.7	66.9	67.8	69.1
<i>ViT-L</i>				
MAE† [26]	42.2	47.7	-	51.7
MSN [1]	-	-	70.1	-
MAGE	60.3	66.1	67.8	69.6
MAGE-C	68.1	71.9	73.0	74.2

and achieves similar performance as MSN [1], which is the state-of-the-art method for self-supervised label-efficient learning. Moreover, the performance of MAGE-C with ViT-L even surpasses the performance of MSN using 13 images per class (1% of ImageNet-1K), even though MSN uses multi-crop augmentation.

Transfer Learning. Another important property of self-supervised representation is its transferability to different datasets. Following the protocol in [19], we evaluate the transfer learning performance of MAGE pre-trained on ImageNet-1K on 8 datasets under a few-shot setting (25 samples per class). Results are shown in Figure 5: we see that MAGE’s superior performance on ImageNet-1K translates to strong performance on other tasks. Since our method operates on quantized semantic tokens instead of raw pixels, it is likely to be more robust to domain shift.

Fine-tuning. Table 4 shows the fine-tuning performance of MAGE and other self-supervised learning meth-

Table 4. Fine-tuning performance on ImageNet-1K. We report the top-1 accuracy and the improvement over training-from-scratch for different methods (other numbers taken from the respective papers). The ViT models trained from scratch on semantic tokens follow the exact same training setting as the ViT models trained from scratch on original image pixels in [26].

Method	ViT-B	ViT-L
scratch on pixels	82.3	82.6
DINO [6]	82.8 (+0.5)	-
MoCo v3 [13]	83.2 (+0.9)	84.1 (+1.5)
BEiT [2]	83.2 (+0.9)	85.2 (+2.6)
MAE [26]	83.6 (+1.3)	85.9 (+3.3)
CAE [11]	83.9 (+1.6)	86.3 (+3.7)
MVP [53]	84.4 (+2.1)	86.3 (+3.7)
PeCo [18]	84.5 (+2.2)	86.5 (+3.9)
scratch on tokens	80.7	80.9
MAGE	82.5 (+1.8)	83.9 (+3.0)
MAGE-C	82.9 (+2.2)	84.3 (+3.4)

ods, when we can change all the pre-trained encoder parameters. Our method achieves performance at par with DINO [5] and slightly under MoCo-v3 [13]. We believe that the use of quantized tokens leads to a subpar from-scratch and fine-tune performance, and leave further investigations of this phenomenon to future work. We note, however, that our method still improves over our supervised training-from-scratch baseline by as large a margin as other methods.

4.4. Analysis

In this section, we analyze the two key components of MAGE that enables the unification of generative modeling and representation learning: variable masking ratio and quantized tokenization. All experiments are conducted on ViT-B. Experiments on variable masking ratio are all trained for 400 epochs, and experiments on quantized tokenization are all trained for 1600 epochs. More analysis and ablation

Table 5. Top-1 accuracy of linear probing and class unconditional generation FID of MAGE on ImageNet-1k with different masking ratio distribution. μ denotes the mode and σ the standard deviation of the truncated Gaussian distribution. When $\sigma = 0$, the masking ratio is fixed and generation has poor quality with very high FID (>50). Therefore we put N/A for FID in such cases.

	$\mu = 0.7$ $\sigma = 0$	$\mu = 0.6$ $\sigma = 0$	$\mu = 0.55$ $\sigma = 0$	$\mu = 0.5$ $\sigma = 0$	$\mu = 0.45$ $\sigma = 0$	$\mu = 0.55$ $\sigma = 0$	$\mu = 0.55$ $\sigma = 0.15$	$\mu = 0.55$ $\sigma = 0.25$	$\mu = 0.55$ $\sigma = 0.5$
Linear Probing	69.7	70.1	71.5	70.9	70.4	71.5	72.0	72.2	71.8
FID	N/A	N/A	N/A	N/A	N/A	N/A	12.5	12.2	13.0

Table 6. Reconstruction loss and linear probe accuracy of MAGE with unquantized features and quantized tokens as input. Using unquantized features makes it much easier to infer masked tokens, and hence results in worse linear probe performance.

inputs	recon. loss	linear probe (%)
Unquantized features	3.31	49.5
Quantized tokens	5.76	74.7

studies are in the Appendix.

Masking Design. Variable masking ratio is one of our key components. We find that the quality of the learned representation can also be affected by the distribution used to sample our masking ratio. We compare the results of MAGE on linear probing and class unconditional generation on ImageNet-1k using different masking ratio distributions in Table 5. We denote the mode of the truncated Gaussian distribution as μ , and the standard deviation of the truncated Gaussian distribution as σ . Note that $\sigma = 0$ indicates a fixed masking ratio. The left 5 columns ablate μ , and the right 4 columns ablate σ . The results show that a variable masking ratio is necessary to enable generation. Moreover, using a variable masking ratio also enables representation learning to learn better features and achieve better linear probe performance.

Tokenization. Previous self-supervised learning methods on images typically directly use raw images as the inputs of the transformer. However, in MAGE we use quantized semantic tokens as both inputs and reconstruction targets. We elaborate the benefits of this design as follows:

- First, during generation, it allows the network to iteratively use its output as the input in the next iteration, which enables high-quality and diverse image reconstruction and generation, as shown in Figure 2 and Figure 4.
- Second, it allows the whole network to operate at a semantic level without losing low-level details and thus extract better representations. We demonstrate this by comparing the linear probe performance on features from each transformer block of ViT-B trained using MAE and MAGE. As shown in Figure 6, the linear probe accuracy of MAGE at each transformer block is always higher than MAE throughout the encoder.
- Third, the quantizer prevents shortcuts created by the VQGAN CNN encoder. If we directly use features extracted




Figure 6. Linear probe accuracy of MAE and MAGE at different transformer blocks of ViT-B. MAGE consistently has higher accuracy across all transformer blocks due to the semantic nature of the quantized tokens.

by the VQGAN encoder without quantization as input to the transformer, since the receptive fields of neighboring feature pixels have significant overlap, it is much easier to infer masked feature pixels using nearby unquantized feature pixels. As shown in Table 6, with the same masking strategy, using unquantized features achieves much lower reconstructive loss (3.31 vs. 5.76), but also a much lower linear probe accuracy (49.5% vs. 74.7%). This suggests that the pre-training task is too easy, leading to shortcut solutions, and hence to poor representations. The quantization step is therefore necessary to learn good representations.

5. Discussion

We have presented MAGE, a masking-based approach that unifies image generation and representation learning in a simple and effective framework. The key to our method is the use of quantized tokens and the use of variable masking ratios that adapt smoothly to both tasks (generation and representation). We have shown extensive results on linear probing, few-shot transfer learning, and unconditional image generation. To the best of our knowledge, this is the first model that achieves close to SOTA results for both tasks using the same data and training paradigm. A natural future extension is to pre-train on larger unlabeled datasets such as JFT300 to further improve performance.

References

- [1] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Michael Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. *arXiv preprint arXiv:2204.07141*, 2022.
- [2] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *Int. Conf. on Learning Representations (ICLR)*, 2019.
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*, 2020.
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Int. Conference on Computer Vision (ICCV)*, pages 9650–9660, 2021.
- [7] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11315–11325, 2022.
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *icml*, pages 1597–1607. PMLR, 2020.
- [10] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in Neural Information Processing Systems*, 33, 2020.
- [11] Xiaokang Chen, Mingyu Ding, Xiaodi Wang, Ying Xin, Shentong Mo, Yunhao Wang, Shumin Han, Ping Luo, Gang Zeng, and Jingdong Wang. Context autoencoder for self-supervised representation learning. *arXiv preprint arXiv:2202.03026*, 2022.
- [12] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*, 2020.
- [13] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Int. Conference on Computer Vision (ICCV)*, pages 9640–9649, 2021.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [15] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.
- [16] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- [17] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. *Advances in neural information processing systems*, 32, 2019.
- [18] Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. Peco: Perceptual codebook for bert pre-training of vision transformers. *arXiv preprint arXiv:2111.12710*, 2021.
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Int. Conf. on Learning Representations (ICLR)*, 2021.
- [21] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Mastropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- [22] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021.
- [23] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- [24] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. 2014.
- [25] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- [26] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16000–16009, June 2022.
- [27] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.

- [28] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- [29] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [30] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 6840–6851, 2020.
- [31] Zhicheng Huang, Xiaojie Jin, Chengze Lu, Qibin Hou, Ming-Ming Cheng, Dongmei Fu, Xiaohui Shen, and Jiashi Feng. Contrastive masked autoencoders are stronger vision learners. *arXiv:2207.13532v1*, 2022.
- [32] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [33] Doyup Lee, Chihoon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [34] Tianhong Li, Peng Cao, Yuan Yuan, Lijie Fan, Yuzhe Yang, Rogerio S Feris, Piotr Indyk, and Dina Katabi. Targeted supervised contrastive learning for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6928, 2022.
- [35] Tianhong Li, Lijie Fan, Yuan Yuan, Hao He, Yonglong Tian, Rogerio Feris, Piotr Indyk, and Dina Katabi. Making contrastive learning robust to shortcuts. *arXiv preprint arXiv:2012.09962*, 2020.
- [36] Hao Liu, Xinghua Jiang, Xin Li, Antai Guo, Deqiang Jiang, and Bo Ren. The devil is in the frequency: Geminated gestalt autoencoder for self-supervised visual pre-training. *arXiv preprint arXiv:2204.08227*, 2022.
- [37] Steven Liu, Tongzhou Wang, David Bau, Jun-Yan Zhu, and Antonio Torralba. Diverse image generation via self-conditioned gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14286–14295, 2020.
- [38] Mario Lučić, Michael Tschannen, Marvin Ritter, Xiaohua Zhai, Olivier Bachem, and Sylvain Gelly. High-fidelity image generation with fewer labels. In *International conference on machine learning*, pages 4183–4192. PMLR, 2019.
- [39] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016.
- [40] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [41] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [42] Zhiliang Peng, Li Dong, Hangbo Bao, Qixiang Ye, and Furu Wei. Beit v2: Masked image modeling with vector-quantized visual tokenizers. *arXiv preprint arXiv:2208.06366*, 2022.
- [43] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [44] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019.
- [45] Joshua Robinson, Li Sun, Ke Yu, Kayhan Batmanghelich, Stefanie Jegelka, and Suvrit Sra. Can contrastive learning avoid shortcut solutions? *Advances in neural information processing systems*, 34:4974–4986, 2021.
- [46] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [47] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [48] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.
- [49] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Int. Conf. on Learning Representations (ICLR)*, 2021.
- [50] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- [51] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [52] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14668–14678, 2022.
- [53] Longhui Wei, Lingxi Xie, Wengang Zhou, Houqiang Li, and Qi Tian. Mvp: Multimodality-guided visual pre-training. *arXiv preprint arXiv:2203.05175*, 2022.
- [54] Yinghao Xu, Yujun Shen, Jiapeng Zhu, Ceyuan Yang, and Bolei Zhou. Generative hierarchical features from synthesizing images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4432–4442, 2021.
- [55] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge,

- and Yonghui Wu. Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627*, 2021.
- [56] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *Int. Conference on Machine Learning (ICML)*, pages 7354–7363, 2019.
- [57] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017.
- [58] Long Zhao, Zizhao Zhang, Ting Chen, Dimitris Metaxas, and Han Zhang. Improved transformer for high-resolution gans. *Advances in Neural Information Processing Systems*, 34:18367–18380, 2021.
- [59] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021.