

Diffusion Models as Masked Autoencoders

Chen Wei^{1,2} Karttikeya Mangalam¹ Po-Yao Huang¹ Yanghao Li¹
 Haoqi Fan¹ Hu Xu¹ Huiyu Wang¹ Cihang Xie³ Alan Yuille² Christoph Feichtenhofer¹
¹FAIR, Meta AI ²Johns Hopkins University ³UC Santa Cruz

Abstract

There has been a longstanding belief that generation can facilitate a true understanding of visual data. In line with this, we revisit generatively pre-training visual representations in light of recent interest in denoising diffusion models. While directly pre-training with diffusion models does not produce strong representations, we condition diffusion models on masked input and formulate diffusion models as masked autoencoders (DiffMAE). Our approach is capable of (i) serving as a strong initialization for downstream recognition tasks, (ii) conducting high-quality image inpainting, and (iii) being effortlessly extended to video where it produces state-of-the-art classification accuracy. We further perform a comprehensive study on the pros and cons of design choices and build connections between diffusion models and masked autoencoders. [Project page](#).

1. Introduction

“What I cannot create, I do not understand.”

— Richard P. Feynman, 1988

For many years, there has been a desire to achieve a deeper understanding of visual data through the process of generation. Early approaches, such as deep belief networks [33] and denoising autoencoders [72], employed generative pre-training to initialize deep networks for downstream recognition tasks. As generative models can *create* new samples by approximating the data distribution, it stands to reason that such modeling should simultaneously arrive at a semantic *understanding* of the raw visual data, as required for recognition tasks, following Feynman.

In line with this philosophy, generative language models, *e.g.*, Generative Pre-trained Transformers or GPTs [3], obtain a broad understanding of language and an immense knowledge base, excelling as both a few-shot learner and a pre-trained base model. Nevertheless, recent explorations in vision generative pre-training fall out of favor. For example, GAN-based BiGAN [15, 16] and auto-regressive iGPT [7] noticeably fall short of their concurrent contrastive algorithms despite of using $10\times$ more param-




Figure 1. **Inference process of DiffMAE**, which iteratively unfolds from random Gaussian noise to the sampled output. During training, the model learns to denoise the input at different noise levels (from *top* row to the *bottom*) and simultaneously performs self-supervised pre-training for downstream recognition.

ters. The challenge stems, in part, from the divergent focus: While recognition models primarily focus on the high-level low-frequency structure of images, generation models must also allocate capacity for low-level high-frequency details [60]. Given this discrepancy, it remains an open question whether and how generative pre-training can effectively compete with other self-supervised algorithms on downstream recognition tasks despite its intuitive appeal.

In recent years, the field of image generation has been dominated by denoising diffusion models [67, 35]. These models employ a straightforward process of iteratively refining noisy samples. As a result, the generated images are of impressively high quality, and even better, can generate an extensive range of diverse samples [50, 59, 64]. In light of this progress, we revisit the potential of generative pre-training in the context of diffusion models.

To begin with, we directly fine-tune a pre-trained diffusion model [13] on ImageNet classification [11]. Despite its strong performance for unconditional image generation, the pre-trained diffusion model only yields a marginal improvement in classification when compared to training the same architecture from scratch (Sec. 3), and is outperformed by concurrent self-supervised pre-training algorithms such as Masked Autoencoders (MAE) [30].

MAE demonstrates strong recognition performance by learning to regress pixels of masked patches given the other visible patches. Inspired by MAE, we incorporate masking into diffusion models and cast Diffusion Models as Masked Autoencoders (DiffMAE). We formulate the masked prediction task as a *conditional* generative objective, *i.e.*, to approximate the pixel distribution of the masked region conditioned on the visible region. We learn models with our diffusion approach, within the MAE framework, introducing no extra training cost. During pre-training, our model is trained to denoise the input at different noise levels and learns a strong representation for *recognition and generation*. We evaluate the pre-trained model by fine-tuning on downstream recognition tasks, as well as image inpainting for which the model generates samples by iteratively unfolding from random Gaussian noise, illustrated in Fig. 1.

The diffusion nature of DiffMAE allows it to generate intricate visual details, *e.g.*, of objects (Fig. 2). In contrast, MAE is known to produce blurry reconstructions that lack high-frequency components. Further, DiffMAE maintains strong performance on image and video recognition tasks.

We make the following observations in this work:

(i) DiffMAE is a strong pre-training approach for fine-tuning on downstream *recognition* tasks, obtaining comparable performance to leading self-supervised learning algorithms that focus solely on recognition. When combined with CLIP [57] features, our DiffMAE is even able to outperform recent work that combines MAE and CLIP.

(ii) DiffMAE is able to *generate* high quality images conditioning on masked input. Specifically, DiffMAE generations outperform leading inpainting methods quantitatively and also appear more semantically meaningful.

(iii) DiffMAE can be extended to the *video* domain effortlessly, providing high-quality inpainting and state-of-the-art recognition accuracy, outperforming recent works.

(iv) We reveal a *connection* between MAE and diffusion models, as MAE effectively performs the first inference step of diffusion. In other words, we believe the success of MAE aligns with the philosophy of generation for recognition.

We further perform a comprehensive empirical study to elucidate the pros and cons of the design choices on downstream recognition and inpainting generation tasks.

2. Related Work

Self-supervised learning aims to learn from unlabeled visual data by a pre-text task that is constructed by image/patch operations (*e.g.*, [14, 51, 75, 79, 26, 4]) and spatial-temporal operations (*e.g.*, [28, 48, 25, 52, 73]). Contrastive methods [19] capitalize on augmentation invariance of images and videos [77, 31, 8, 5, 29, 6, 9, 56, 24].

For vision, different masked prediction targets have been proposed. MAE [30] predicts pixel colors with an efficient asymmetric architecture. BEiT [2, 55] and iBOT [82] pre-

dict dVAE [69, 60] or learnable tokens. MaskFeat [74] predicts HOG features. data2vec [1] learns from a momentum teacher. The community is continuously exploring this direction [17, 78, 38, 81, 46, 22].

Generative learning for recognition has a long-standing appeal for its intuitiveness. Pioneers study the representation learned by GANs [27, 58, 15] and VAEs [40, 39]. Big-BiGAN [16] demonstrates a model that learns competitive recognition representation and generates high-fidelity images with GAN. This work is followed by iGPT [7], which generates images autoregressively and was state-of-the-art on linear probing protocols. Though this stream stagnates afterward, we demonstrate its resurgence by exploring diffusion models in this work.

Denoising diffusion models have ushered in a new era of diverse, high-resolution conditional image generation [67, 13, 35]. Utilizing a forward Gaussian diffusion process and a backward generation process, denoising diffusion models iteratively refine the generated image starting from Gaussian noise. This process has proven extremely powerful for rich text-conditioned generation of both images [59, 50, 64, 62] and videos [71, 36, 34, 66].

Masked autoencoders were pioneered by stacked autoencoders [72] and inpainting tasks [53] using ConvNets. Since the introduction of ViT [18], masked prediction has re-attracted attention, partially inspired by the success of BERT [12] in NLP. BERT performs masked language modeling, scales well, and generalizes to different end tasks.

3. Fine-tuning Generative Models

pre-train	architecture	params.	scratch	pre-trained
MAE [30]	ViT-L	304M	82.6	85.9
iGPT [7]	iGPT-L	1362M	53.2	72.6
ADM [13]	U-Net enc.	211M	82.0 [†]	83.3
DDPM [35]	ViT-L	304M	82.6	83.4
DiffMAE	ViT-L	304M	82.6	85.8

Table 1. **Fine-tuning generative models on ImageNet**, a system-level comparison. [†]trained from scratch by us. Non-generative method is included for reference and *de-emphasized*.

We start by comparing different generative pre-training methods for downstream ImageNet classification in Tab. 1.

iGPT [7] generates images in a GPT [3] style by autoregressively predicting pixels. iGPT pre-training has significantly improved its low from-scratch (random initialization) accuracy from 53.2% to 72.6%, which is still relatively low, considering the model size (1362M) and compute cost.

To re-examine generative pre-training in light of the recent progress on diffusion models, we fine-tune a pre-trained ADM [13], which is a recent method for diffusion-based image generation. In particular, we take its unconditional IN-1K model, which aligns with the criteria of unsupervised pre-training and never uses the class labels.




Figure 2. **Qualitative comparison of DiffMAE and MAE.** $t = 1000$ and $t = 500$ are DiffMAE’s predictions at intermediate timesteps from $t = 1000 \dots 0$. $t = 0$ marks the final generation. MAE’s outputs are obtained from its official visualization model. The predictions at the masked area and the ground-truth at the visible area are overlaid. Both models are ViT-L pre-trained for 1600 epochs.

Compared to the from-scratch counterpart, initializing with the pre-trained diffusion model provides a gain of +1.3% top-1. However, the resulting 83.4% still lags far behind the non-generative self-supervised algorithms such as MAE (85.9%). Considering the differences in architecture, we further train a ViT-L with diffusion DDPM [35]. This model enhances the fine-tuning classification to 83.4% compared to its from-scratch baseline of 82.6% top-1 [30], which aligns with our observations on fine-tuning the pre-trained ADM. Implementation details are in the Appendix.

In comparison to these approaches, a DiffMAE trained ViT-L obtains 85.8% when fine-tuned and is described next.

4. Diffusion Masked Autoencoder

We incorporate the masked prediction paradigm into diffusion models. The model approximates the pixel distribution of the masked region conditioned on the visible region. In other words, we study Diffusion Models as a form of Masked Autoencoders (DiffMAE), introduced next.

4.1. Conditional Diffusion Model

Given a training sample $\mathbf{x}_0 \sim p(\mathbf{x}_0)$ where subscript $_0$ denotes that the sample is original and clean, we first spatially divide \mathbf{x}_0 into the non-overlapping masked region \mathbf{x}_0^m and the visible region \mathbf{x}_0^v . We want to model the distribution of the masked region \mathbf{x}_0^m conditioned on the visible region \mathbf{x}_0^v , *i.e.*, the distribution $p(\mathbf{x}_0^m | \mathbf{x}_0^v)$.

In the forward process, only the masked area \mathbf{x}_0^m is gradually diffused, *i.e.*, corrupted by recursively adding a small amount of Gaussian noise T times to $\mathbf{x}_1^m, \mathbf{x}_2^m, \dots, \mathbf{x}_T^m$ fol-

lowing the Markov process below:

$$p(\mathbf{x}_t^m | \mathbf{x}_{t-1}^m) = \mathcal{N}(\mathbf{x}_t^m; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}^m, \beta_t \mathbf{I}), \quad (1)$$

where $t \in [1, 2, \dots, T]$ denotes the timestep and $\beta_{1:T}$ is the variance schedule of noise, held as hyper-parameters.

Thanks to the properties of Gaussian distribution, we can directly sample \mathbf{x}_t^m without the recursive formulation by:

$$p(\mathbf{x}_t^m | \mathbf{x}_0^m) = \mathcal{N}(\mathbf{x}_t^m; \sqrt{\bar{\alpha}_t} \mathbf{x}_0^m, (1 - \bar{\alpha}_t) \mathbf{I}), \quad (2)$$

which can be reparameterized to $\mathbf{x}_t^m = \sqrt{\bar{\alpha}_t} \mathbf{x}_0^m + \sqrt{1 - \bar{\alpha}_t} \epsilon$ for $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. The variance schedule makes sure that $\bar{\alpha}_T$ of the final timestep T is sufficiently small so that $p(\mathbf{x}_T^m)$, the end of the forward process, approximates the standard normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ well, which is the starting point of the reverse process.

We want to generate masked regions by sampling from $p(\mathbf{x}_0^m | \mathbf{x}_0^v)$, which is approximated by recursively sampling from $p(\mathbf{x}_{t-1}^m | \mathbf{x}_t^m, \mathbf{x}_0^v)$ starting from $\mathbf{x}_T^m \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. When the variance of noise β_t in each step t is small enough, $p(\mathbf{x}_{t-1}^m | \mathbf{x}_t^m, \mathbf{x}_0^v)$ can also be also considered Gaussian distributed [67], to be approximated by a deep network. We optimize the *simple* objective proposed by DDPM [35]:

$$\mathcal{L}_{\text{simple}} = \mathbb{E}_{t, \mathbf{x}_0, \epsilon} \| \mathbf{x}_0^m - D_\theta(\mathbf{x}_t^m, t, E_\phi(\mathbf{x}_0^v)) \|^2. \quad (3)$$

This mean squared error objective is simplified from the variational bound. Unlike DDPM using ϵ -prediction, which predicts the noise, our model predicts the denoised masked region \mathbf{x}_0^m . These two formulations are interchangeable, and both are commonly adopted in diffusion models [59]. E_ϕ is the encoder, which projects the visible image content \mathbf{x}_0^v into the latent space. D_θ is the decoder predicting denoised input, from the noisy \mathbf{x}_t^m , the timestep t , and the visible latent $E_\phi(\mathbf{x}_0^v)$.

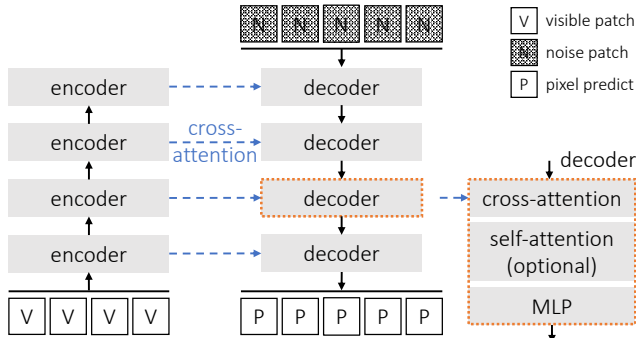


Figure 3. **DiffMAE with cross-self and cross decoder.** The decoder inputs are noisy patches at randomly sampled noise levels. The decoder outputs are the predicted pixels. The decoder blocks are top-down connected to the encoder blocks. Cross-self decoder has self-attention in each block, while cross decoder does not.

4.2. Architecture Design

We instantiate the above conditional diffusion model following the asymmetric design of MAE, introducing no extra training cost. The overall architecture is built solely on Vision Transformers (ViT) [18], which is in contrast to the typical choice of U-Net [63] backbones in diffusion models [35, 13, 59, 64]. In this way, we enable a straightforward evaluation of DiffMAE’s ability on downstream recognition tasks and an apple-to-apple comparison to other self-supervised pre-training methods.

Encoder. The encoder takes a standard ViT. Specifically, the training images are first divided into non-overlapping patches, among which most are selected as the visible patches x_0^v and the others are the masked x_0^m . The ViT encoder $E_\phi(\cdot)$ only operates on the visible patches and encodes each of these patches into the latent space. The encoded $E_\phi(x_0^v)$ then serves as the condition of the generation task performed by the decoder, providing hints of the masked object. After the pre-training stage, only the encoder is fine-tuned to downstream tasks.

Decoder. The decoder takes in noisy masked patches x_t^m as the input. The noise levels of these patches, denoted by the timestep t , are integers uniformly sampled in $[1, T]$ during training. As in ViT, we first project these noisy patches to noise tokens using a linear layer. The timestep t can be specified by adding sinusoidal embeddings to the noise tokens together with the positional embeddings. However, our experiments show that the addition of the t embedding or the lack thereof does not make a big difference to both downstream recognition and inpainting generation, which suggests that the decoder can automatically determine the level t of noisy patches when conditioned on clean patches.

We explore three different decoder configurations, which differ in how the attention modules are applied to the visible latents and the noise tokens:

- *Joint decoder* applies Transformer blocks on the concatenated sequence of the visible latents from the last encoder block and the noise tokens. Each Transformer block operates equally on the visible latents and the noise tokens, with a self-attention layer followed by an MLP.
- *Cross-self decoder* resembles the encoder-decoder design in the original Transformer [70]. In each decoder block, the noise tokens first attend to the visible latents with a *cross-attention* layer and then attend to each other noise tokens with a *self-attention* layer. The cross-attention enables the decoder to attend to the visible latents of different encoder blocks, rather than solely to the final block as in joint decoder. In this way, the encoder and the decoder can be connected in a U-shaped manner illustrated in Fig. 3, which is typically advantageous for image generation and dense predictions such as U-Net [63].
- *Cross decoder* is similar to cross-self decoder, but excludes the self-attentions between the noise tokens. In other words, each noise token independently attends to the visible latents from the encoder in a top-down manner as in Fig. 3, *not* being aware of the existence of other noise tokens. With a much smaller attention map, cross decoder is the most efficient among the three. Specifically, training with cross decoder is overall $\sim 15\%$ more efficient than joint decoder, using a ViT-L encoder and an eight-block decoder of width 512.

We ablate the three decoder architectures in Sec. 5.

Sampling. For inference, the encoder $E_\phi(x_0^v)$ forwards the visible patches *only once* while the decoder unfolds from Gaussian noise to the sampled image for T times iteratively following DDPM [35]. As illustrated in Fig. 2, the predictions gradually become less blurry as the iterative unfolding progresses, and both small-scale structures of the objects (e.g., the mouth of the monkey) and high-frequency details (e.g., the fur of the fox) emerge in the generated images. After T iterations, we obtain high-quality samples at $t = 0$.

CLIP target. To compare to recent prior works that use CLIP [57], we also explore a version for predicting CLIP features together with the original pixel prediction task by simply using a second decoder. The prediction is optimized by minimizing the cosine distance to the CLIP features of the masked patches, similar to MaskFeat [74] and MILAN [37]. As we show in Tab. 6, predicting CLIP features not only enhances recognition ability but also brings improved inpainting generation quality.

Video. The above model are described in the context of *image* operations. However, it can be straightforwardly extended to the spatiotemporal domain of *videos*. Most operations are shared, except that there is an additional temporal dimension. Each token now represents a space-time cube corrupted by noise. The masking is performed by randomly

sampling space-time cubes. The prediction target are pixels of a single time slice of each masked space-time cube.

4.3. Connection between Diffusion and MAE

We discuss the connection between diffusion models and MAE. The two are different at first glance: They have different purpose (image generation vs. self-supervised pre-training), inputs (noisy images vs. masked images), outputs (noise vs. pixels) and architectures (U-Net vs. ViT). Consider Fig. 2, at the first timestep $t = 1000$ where the inputs of the decoder approximate Gaussian noise, the pixel predictions of our model are visually similar to the predictions of MAE, both capturing only the blurry and rough structures of the objects. Moreover, DiffMAE trained with only the $t = 1000$ noise level obtains a similar fine-tuning accuracy of around 85.0% as MAE, shown in Fig. 4. This consistently observed similarity makes sense because both the Gaussian noise in DiffMAE at $t = 1000$ and the learnable mask token in MAE contain no image signal.

This observation suggests a close relationship between MAE and diffusion models, that MAE effectively performs the first inference step ($t = 1000$) of diffusion models. On the other hand, the other steps of diffusion models generate more and more intricate high-frequency details, which is lacking in the predictions of MAE. Overall, the connection suggests that MAE can be viewed as a single-step patch-conditioned diffusion model, and that the success of MAE in downstream recognition tasks is *in line with* the philosophy of generative pre-training. Correspondingly, diffusion models, the first step of whom effectively performs MAE, are potentially good recognition models.

5. Empirical Study

Settings. We pre-train our DiffMAE on the IN-1K [11] training set. If not specified, the encoder is a vanilla ViT-L [18] without any modification, and the decoder is of depth 8 and width 512. The eight cross-(self) decoder blocks spread top-down uniformly attending to the outputs of the 24 encoder blocks of ViT-L. The data augmentation is randomly resized cropping, and the masking strategy is to mask out 75% patches randomly. All models are pre-trained for 400 epochs. For the diffusion setup, the variance follows a linear schedule [35], and the number of timesteps T is set to 1000. Details are in the Appendix.

From the perspective of pre-training for *recognition*, we report end-to-end fine-tuning top-1 accuracy on the IN-1K validation set. The model is both pre-trained and fine-tuned at 224^2 resolution. By default, we use our cross decoder design, and the prediction targets are per-patch normalized.

For the task *generative* inpainting, we report the common LPIPS \downarrow [80], which is a learned distance metric based on the deep feature space of AlexNet [41]. Following the

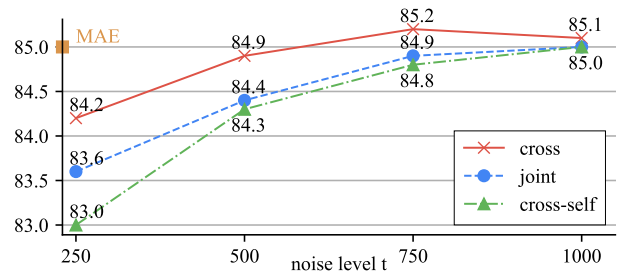


Figure 4. **Comparison of different decoder architectures** across noise levels (x -axis) by fine-tuning top-1 (%) on IN-1K val set (y -axis). Each DiffMAE model is pre-trained with a *single* noise level. All models are ViT-L pre-trained for 400 epochs.

setting of RePaint [47], we compute LPIPS \downarrow over 100 images from the IN-1K validation set, and train and evaluate the model with image size 256^2 . We use cross-self decoder, and the prediction targets are unnormalized pixels.

Decoder architecture – classification. We find that the three decoder architectures impact fine-tuning for downstream classification differently across different noise levels. To demonstrate this, we train the models at various *single* noise levels, analyzed in Fig. 4. A smaller t corresponds to less noise while a larger t denotes more noise. At $t = 1000$, the inputs have almost no image signals but noise. At $t = 0$, the inputs would be the clean image patches.

First, the fine-tuning accuracy drops for all three decoders when t is lower. At a small noise level t , the disruption to image signals by noise is less severe. Consequently, the decoder needs less conditioning on the visible latents to predict image signals, which weakens the pre-training of the encoder. On the other hand, the model is now tasked to generate more low-level, high-frequency details, which are considered to be less useful for downstream recognition.

Second, the drop rates of joint and cross-self decoder are faster than that of cross decoder. This is because both joint decoder and cross-self decoder have a shortcut to exploit the image signals of the neighboring noisy patches and bypass the visible latents of the encoder. This shortcut is amplified especially when t is small. In contrast, cross decoder, whose inputs do not see other noisy patches, avoids this shortcut.

In summary, our experiments show that the pre-training performance is closely related to the difficulty of the denoising task. Nonetheless, the pre-training helps recognition in all cases compared to its from-scratch counterpart of 82.6% top-1 [30]. We make cross decoder the default option for pre-training for downstream recognition tasks.

Decoder architecture – inpainting. Next, we discuss the effect of different decoders on inpainting. For the cross decoder, the noise tokens are independent of each other and do not see other noise tokens. Therefore, the generated image is patch-wise discontinuous. Even so, these patches still make sense semantically at their positions, show in Fig. 5.




Figure 5. **Generations of cross and cross-self decoders.** The generation of cross decoder is patch-wise discontinuous, while cross-self decoder generates visually continuous samples.

The generations of both joint and cross-self decoders are continuous and visually pleasant. Comparing these two, cross-self decoder achieves a better LPIPS \downarrow score of 0.216 on the evaluation set while joint decoder obtains 0.224. We credit the better generation of cross-self decoder to its U-shape design following U-Net, which is commonly suitable for dense predictions. We make cross-self decoder the default option for inpainting.

range of t	ρ	top-1
1000	1.0	85.1
0~1000	1.0	84.7
250~1000	1.0	84.7
500~1000	1.0	85.1
0~1000	0.8	85.1
0~1000	0.6	85.0

Table 2. **Noise schedule.** We first vary the range of t with different starting points. We then modify the default linear schedule with a hyper-parameter ρ as the exponent to each variance β_t .

Noise variance schedule. Next, we train our model with mixed noise levels to complete the diffusion process in Tab. 2. We use the linear schedule following DDPM [35] by default, where the variances $\beta_{[1:T]}$ increase linearly.

We first vary the range of t , starting from $t = 250$ and $t = 500$. Similar to our observations on training with single noise levels, the less the t range includes lightly noisy samples, the better the fine-tuning accuracy is. When t is mixed from 500 to 1000, the pre-trained model is just as good as the case of training with only $t = 1000$ samples.

We then directly modify the linear schedule with a hyper-parameter ρ , which exponentiates the variances β_t to β_t^ρ . When ρ is smaller than one, it enlarges each variance β_t , therefore amplifying the amount of noise at each timestep. Consequently, the fine-tuning top-1 improves with a smaller ρ from 84.7% to around 85.0%. However, the gain on downstream classification is at the cost of a drop in the inpainting generation. Specifically, LPIPS \downarrow adversely increases from 0.216 with the default linear schedule ($\rho = 1.0$) to 0.228 with $\rho = 0.6$. This suggests another trade-off between pre-training for recognition and inpainting generation. We detail ρ formulation in the Appendix.

Prediction target. We next study the influence of prediction target in Tab. 3. Both noise and pixel prediction are commonly adopted in generative diffusion models [35, 59].

target	noise	pixel	pixel w/ norm
top-1	unstable	84.3	85.1

Table 3. **Prediction targets** of noise, pixels, and per-patch normalized pixels. The noise entry is unstable in fine-tuning.

However, when evaluated for downstream classification, only predicting the pixels can serve as the target, obtaining 84.3% top-1. The noise prediction, however, can not be stably fine-tuned for the downstream classification task. Regarding the pixel prediction, the per-patch normalization [30] consistently helps the fine-tuning accuracy.

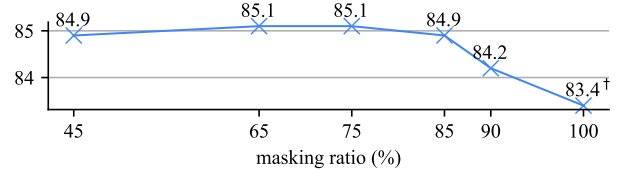


Figure 6. **Masking ratio.** Varying the masking ratio of random patch masking. [†]For 100% masking, we use the encoder-only architecture instead of the asymmetric encoder-decoder architecture.

Masking ratio. We consider different masking ratios for the random masking strategy. We report the fine-tuning accuracy in Fig. 6. We observe that accuracy is stable with masking ratios from 45% to 85% and has a drop when the masking ratio increases to 90%. When the masking ratio is increased to 100%, *i.e.*, no masking at all, the fine-tuning top-1 drops to 83.4%. This suggests the importance of visible patch conditioning for DiffMAE pre-training. On the other hand, the recognition performance is still enhanced compared to its from-scratch baseline of 82.6% top-1 [30].

mask	random	center	training	masking	random	center
top-1	85.1	84.3	center	LPIPS \downarrow	0.142	0.125

Table 4. **Center masking.** The left reports the fine-tuning top-1. The right reports LPIPS \downarrow of the two models trained with random and center mask, respectively, but evaluated *both with center mask*.

Center masking. We now compare random masking with a more common evaluation protocol for inpainting tasks, center masking, which block-wise masks out the center 25% region of images as in Fig. 7. We first train with two different maskings and then evaluate both models by fine-tuning top-1. As shown in Tab. 4 (*left*), random 75% mask works better for fine-tuning.

We then evaluate *both random and center mask* trained models by inpainting LPIPS \downarrow with *center mask* in Tab. 4 (*right*), the common inpainting protocol. The random mask trained model provides reasonable center block inpainting, reporting 0.142 LPIPS \downarrow on the evaluation set. This shows that the model still obtains block-wise inpainting ability although trained with random masks. When trained with center mask, the inpainting quality is further improved, and LPIPS \downarrow favorably decreases to 0.125.




Figure 7. **Samples of center masking inpainting** of DiffMAE. For a 256^2 image, we mask out the center 128^2 block.

depth	top-1	LPIPS \downarrow	width	top-1	LPIPS \downarrow
4	85.0	0.220	256	85.0	0.239
8	85.1	0.216	512	85.1	0.216
12	85.1	0.215	1024	84.9	0.208

Table 5. **Decoder size.** The left table varies the depth of the decoder with the default width of 512. The right table varies the width of the decoder with the default depth of eight blocks.

Decoder size. We discuss the effect of the decoder size in Tab. 5. By default, we use eight decoder blocks, each of which has 512 channels. We do not observe significant differences in the fine-tuning performance when varying the decoder depth or the width, all of which have around 85.0% top-1. For inpainting, we observe a notable drop if the decoder size is reduced and a gain when the decoder size is enlarged. With only four decoder blocks, the LPIPS \downarrow adversely increases from 0.216 to 0.220. A narrower decoder of width 256 leads to an even worse LPIPS \downarrow of 0.239. Using an extra wide decoder with 1024 channels, the inpainting is largely improved to 0.208 LPIPS \downarrow .

target	pixel	CLIP	pixel+CLIP	target	pixel	pixel+CLIP
top-1	85.1	86.5	86.7	LPIPS \downarrow	0.226	0.216

Table 6. **CLIP target.** Varying the prediction target between pixel, CLIP feature, and the combination of the two.

CLIP target. Boosting pre-training with CLIP [57], a vision-language model trained on a 400M in-house dataset, has been popular driven by its strong performance in downstream tasks. MaskFeat [74] first showed improved performance for masked feature prediction. Then algorithms utilizing CLIP either regress CLIP features [76, 37], or predict CLIP-aided tokens [55]. In this work, we use a second cross decoder to regress CLIP features alongside the original pixel decoder, multitasking on both masked pixel and CLIP feature prediction. In Tab. 6 (left), we first compare the options by fine-tuning. Solely regressing CLIP features achieves a high accuracy 86.5%, compared to the default pixel predicting baseline 85.1%. Multitasking on both targets further improves the accuracy to 86.7%.

Beyond the improvement on fine-tuning, we observe in Tab. 6 (right) that combining CLIP targets can also improve inpainting. The LPIPS \downarrow decreases from 0.226 to 0.216 when multitasking with CLIP regression. This suggests that a better semantic understanding, in this case by learning from CLIP, can notably help image inpainting quality.

pre-train	w/ CLIP	ViT-B	ViT-L	ViT-H
from-scratch [30]	×	82.3	82.6	83.1
MoCo v3 [10]	×	83.2	84.1	-
DINO [6]	×	82.8	-	-
iBOT [82]	×	84.0	84.8	-
BEiT [2]	×	83.2	85.2	-
MaskFeat [74]	×	84.0	85.7	-
MAE [30]	×	83.6	85.9	86.9
DiffMAE	×	83.9	85.8	86.9
MVP [76]	✓	84.4	86.3	-
MILAN [37]	✓	85.4	86.7	-
BEiTv2 [†] [55]	✓	85.5	87.3	-
DiffMAE	✓	84.9	86.9	88.0


Table 7. **Comparison to prior works on IN-1K fine-tuning.** All entries are pre-trained on IN-1K train split at image size 224^2 . When using CLIP, all ViT-B and ViT-L utilize CLIP ViT-B/16 and ViT-H uses CLIP ViT-L/14. [†]BEiTv2’s ViT use layer scale [68] and relative positional embedding [65].

Discussion. We have thoroughly studied the design aspects of DiffMAE in terms of downstream classification and generative inpainting performance. Our experiments show that, for many aspects, pre-training for downstream classification and generative inpainting do not share the optimal setting, *e.g.*, the decoder architecture and the noise variance schedule. For other aspects, settings favoring inpainting quality do not impact pre-training ability, *e.g.*, decoder size. Interestingly, using CLIP benefits both downstream classification and generative inpainting. We use cross decoder, per-patch normalization, and noisy variance schedule for fine-tuning evaluation, and we use cross-self decoder with 1024 width, the default linear schedule, but not per-patch normalization for inpainting evaluation. While the inpainting-oriented model can still improve from-scratch recognition accuracy, we still would find it desirable to have an identical setting that is optimal for both tasks.

6. Comparison to Prior Work

ImageNet fine-tuning. We compare DiffMAE to prior works in fine-tuning for IN-1K classification in Tab. 7.

The first section presents methods that are solely trained on IN-1K training set without access to CLIP models. We pre-train DiffMAE for 1600 epochs here while 400 were used for ablations in Sec. 5. We observe that fine-tuning DiffMAE significantly boosts the scratch accuracy for all ViT-B (+1.6%), ViT-L (+3.2%) and ViT-H (+3.8%) models. This suggests that DiffMAE can provide strong representations for recognition compared to prior generative models (*cf.* Tab. 1), and is scalable to large models. Compared to prior self-supervised learning works, it outperforms contrastive methods like MoCo v3 and DINO and is comparable masked autoencoders, while being able to *generatively inpaint images with high quality* at the same time.



original 75% random RePaint MAE DiffMAE
 Figure 8. **Visualizations on inpainting** of three different algorithms. Generations of DiffMAE have finer details than MAE and are more semantically meaningful than RePaint.

The second section presents representative works that use CLIP to enhance representation learning. We pre-train our DiffMAE for 800 epochs multitasking in masked pixel and CLIP feature prediction. For ViT-B and ViT-L models, all entries utilize CLIP ViT-B/16. For ViT-H, we use CLIP ViT-L/14. The CLIP models are pre-trained on in-house WIT400M [57] and lift accuracy for both DiffMAE and other works. Combined with CLIP, DiffMAE obtains a strong IN-1K top-1 of 88% with ViT-H, while ViT-B lags behind because of its insufficient capacity for multitasking.

arch.	pre-train	AP ^{box}	AP ^{mask}
Swin-L	Sup., IN-21K	52.4	46.2
ViT-L	MAE, IN-1K	55.6	49.2
ViT-L	DiffMAE, IN-1K	55.3	49.0

Table 8. **Comparisons on COCO instance segmentation.**

COCO instance segmentation. We conduct COCO [44] instance segmentation with ViTDet [42] using Mask R-CNN [32] in Tab. 8. Our results are pre-trained on IN-1K and slightly lower than those of MAE, while being much stronger than *e.g.* supervised Swin-L trained on IN-21K.

Video classification. We compare DiffMAE to prior work on K400 fine-tuning for video classification in Tab. 9.

The first section presents results that use no in-house data. Our model largely improves from the from-scratch models and achieves comparable results to other algorithms, which use masked autoencoding methodology for pre-training. In the second section, we present the results that use extra in-house data. Especially, we compare to EVL [45] and X-CLIP [49], both using CLIP ViT-L/14 model, which is trained on the in-house text-image dataset WIT400M [57]. DiffMAE obtains 87.4% top-1 with a regular ViT-L/14 model, outperforming these works. This number is further improved to a high 88.1% with a larger input size 280² and a longer temporal duration of 32 frames.

ImageNet inpainting. We first qualitatively compare RePaint [47], MAE, and our DiffMAE on 75% random mask inpainting in Fig. 8. RePaint employs a pre-trained unconditional diffusion model, specifically ADM [13], as the gener-

model	pre-train	extra data	top-1	top-5	input size	FLOPs	param
SlowFast [23]	scratch	-	79.8	93.9	64×224 ²	234×3×10	60
X3D-XL [21]	scratch	-	79.1	93.9	16×312 ²	48×3×10	11
MViT-B [20]	scratch	-	81.2	95.1	64×224 ²	455×3×3	37
MViTv2-B [43]	scratch	-	82.9	95.7	32×224 ²	255×1×5	51
MViTv2-L [43]	MaskFeat [74]	-	84.3	96.3	16×224 ²	377×1×10	218
ViT-L [18]	MAE [22]	-	84.8	96.2	16×224 ²	598×3×7	304
ViT-L [18]	DiffMAE	-	84.5	96.3	16×224 ²	598×3×7	304
ViT-L/14 [18]	EVL [45]	WIT400M	86.9	97.4	16×224 ²	1348×1×3	n/a
ViT-L/14 [18]	X-CLIP [49]	WIT400M	87.1	97.6	8×224 ²	658×4×3	n/a
ViT-L/14 [18]	DiffMAE	WIT400M	87.4	97.5	16×224 ²	828×3×10	304
ViT-L/14 [18]	EVL [45]	WIT400M	87.7	97.8	32×336 ²	6065×1×3	n/a
ViT-L/14 [18]	X-CLIP [49]	WIT400M	87.7	97.4	16×336 ²	3086×4×3	n/a
ViT-L/14 [18]	DiffMAE	WIT400M	88.1	97.8	32×280 ²	2588×4×3	304

Table 9. **Comparison with previous works on Kinetics-400.** We report the inference cost with a single “view” (temporal clip with spatial crop) × the number of views (FLOPs × view_{space} × view_{time}). Magnitudes are Giga (10⁹) for FLOPs and Mega (10⁶) for Param.

method	arch.	random	center
DSI [54]	VQ-VAE2 [61]	0.300	0.153
RePaint [47]	ADM [13]	0.303	0.160
DiffMAE	ViT-L	0.208	0.125
DiffMAE	ViT-H	0.205	0.121

Table 10. **Comparison to prior works on IN-1K inpainting** with LPIPS_↓ for 75% random mask and 25% center mask settings. All entries are trained and evaluated at image size 256².

ative prior to perform free-form inpainting. We notice that the generations of RePaint, though with visually pleasant details, often need more precise semantic meaning. In contrast, MAE generates blurry samples. DiffMAE generates visually detailed and semantically meaningful images.

We finally quantitatively compare the inpainting performance in Tab. 10. Both DSI and RePaint are algorithms specially designed for inpainting. We evaluate at image size 256² following both works. We first compare the algorithms in the 75% random mask setting where DSI and RePaint’s results are obtained by us with their official models. We then compare the three algorithms with the 25% center mask protocol, which is commonly used by inpainting algorithms. DiffMAE largely improves LPIPS_↓ on both random 75% mask and center block mask protocols, setting a new state-of-the-art for ImageNet inpainting. We include more generative inpainting examples in the Appendix.

7. Conclusion

We present Diffusion Masked Autoencoders (DiffMAE), a self-supervised framework designed for recognizing and generating images and videos. We accomplish this by integrating masking into diffusion models, thereby transforming conditional diffusion models into masked autoencoders. By re-evaluating the philosophy of generation and viewing it as a tool for genuinely comprehending visual data, we contribute to the current trend in generative models. Ultimately, we aspire for our research to spark further exploration of integrated methods for recognition and generation.

Acknowledgement. We thank Yuan-Ting Hu and Minkai Xu for valuable discussion. CW and AY is supported by ONR N00014-21-1-2690.

References

- [1] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. In *ICML*, 2022. 2
- [2] Hangbo Bao, Li Dong, and Furu Wei. BEiT: BERT pre-training of image transformers. In *ICLR*, 2022. 2, 7
- [3] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *NeurIPS*, 2020. 1, 2
- [4] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 2
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 2
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 2, 7
- [7] Mark Chen, Alec Radford, Rewon Child, Jeff Wu, Heewoo Jun, Prafulla Dhariwal, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *ICML*, 2020. 1, 2
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2
- [9] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, 2021. 2
- [10] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021. 7
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 1, 5
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 2
- [13] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 1, 2, 4, 8
- [14] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 2
- [15] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. In *ICLR*, 2017. 1, 2
- [16] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. In *NeurIPS*, 2019. 1, 2
- [17] 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. 2
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2, 4, 5, 8
- [19] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *TPAMI*, 2015. 2
- [20] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. 8
- [21] Christoph Feichtenhofer. X3D: Expanding architectures for efficient video recognition. In *CVPR*, 2020. 8
- [22] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. In *NeurIPS*, 2022. 2, 8
- [23] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 8
- [24] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross Girshick, and Kaiming He. A large-scale study on unsupervised spatiotemporal representation learning. In *CVPR*, 2021. 2
- [25] Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *CVPR*, 2017. 2
- [26] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 2
- [27] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *NeurIPS*, 2014. 2
- [28] Ross Goroshin, Joan Bruna, Jonathan Tompson, David Eigen, and Yann LeCun. Unsupervised learning of spatiotemporally coherent metrics. In *ICCV*, 2015. 2
- [29] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning. In *NeurIPS*, 2020. 2
- [30] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 1, 2, 3, 5, 6, 7
- [31] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 2
- [32] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 8
- [33] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 2006. 1

- [34] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 2
- [35] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 1, 2, 3, 4, 5, 6
- [36] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *NeurIPS*, 2022. 2
- [37] Zejiang Hou, Fei Sun, Yen-Kuang Chen, Yuan Xie, and Sun-Yuan Kung. MILAN: Masked image pretraining on language assisted representation. *arXiv preprint arXiv:2208.06049*, 2022. 4, 7
- [38] Ronghang Hu, Shoubhik Debnath, Saining Xie, and Xinlei Chen. Exploring long-sequence masked autoencoders. *arXiv preprint arXiv:2210.07224*, 2022. 2
- [39] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. *NeurIPS*, 2014. 2
- [40] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 2
- [41] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 5
- [42] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *ECCV*, 2022. 8
- [43] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Improved multiscale vision transformers for classification and detection. In *CVPR*, 2022. 8
- [44] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 8
- [45] Ziyi Lin, Shijie Geng, Renrui Zhang, Peng Gao, Gerard de Melo, Xiaogang Wang, Jifeng Dai, Yu Qiao, and Hongsheng Li. Frozen clip models are efficient video learners. In *ECCV*, 2022. 8
- [46] Xingbin Liu, Jinghao Zhou, Tao Kong, Xianming Lin, and Rongrong Ji. Exploring target representations for masked autoencoders. *arXiv preprint arXiv:2209.03917*, 2022. 2
- [47] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. RePaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, 2022. 5, 8
- [48] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016. 2
- [49] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding language-image pretrained models for general video recognition. In *ECCV*, 2022. 8
- [50] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *ICML*, 2022. 1, 2
- [51] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 2
- [52] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *CVPR*, 2017. 2
- [53] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 2
- [54] Jialun Peng, Dong Liu, Songcen Xu, and Houqiang Li. Generating diverse structure for image inpainting with hierarchical vq-vae. In *CVPR*, 2021. 8
- [55] 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. 2, 7
- [56] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. In *CVPR*, 2021. 2
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 4, 7, 8
- [58] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016. 2
- [59] 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. 1, 2, 3, 4, 6
- [60] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. 1, 2
- [61] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with VQ-VAE-2. In *NeurIPS*, 2019. 8
- [62] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [63] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 4
- [64] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022. 1, 2, 4
- [65] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. In *NAACL*, 2018. 7

- [66] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. [2](#)
- [67] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. [1](#), [2](#), [3](#)
- [68] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *ICCV*, 2021. [7](#)
- [69] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *NeurIPS*, 2017. [2](#)
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [4](#)
- [71] Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable length video generation from open domain textual description. *arXiv preprint arXiv:2210.02399*, 2022. [2](#)
- [72] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *JMLR*, 2010. [1](#), [2](#)
- [73] Xiaolong Wang, Kaiming He, and Abhinav Gupta. Transitive invariance for self-supervised visual representation learning. In *ICCV*, 2017. [2](#)
- [74] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *CVPR*, 2022. [2](#), [4](#), [7](#), [8](#)
- [75] Chen Wei, Lingxi Xie, Xutong Ren, Yingda Xia, Chi Su, Jiaying Liu, Qi Tian, and Alan L Yuille. Iterative reorganization with weak spatial constraints: Solving arbitrary jigsaw puzzles for unsupervised representation learning. In *CVPR*, 2019. [2](#)
- [76] Longhui Wei, Lingxi Xie, Wengang Zhou, Houqiang Li, and Qi Tian. MVP: Multimodality-guided visual pre-training. In *ECCV*, 2022. [7](#)
- [77] Zhirong Wu, Yuanjun Xiong, X Yu Stella, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. [2](#)
- [78] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. SimMIM: A simple framework for masked image modeling. In *CVPR*, 2022. [2](#)
- [79] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016. [2](#)
- [80] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [5](#)
- [81] Xiaosong Zhang, Yunjie Tian, Wei Huang, Qixiang Ye, Qi Dai, Lingxi Xie, and Qi Tian. HiViT: Hierarchical vision transformer meets masked image modeling. *arXiv preprint arXiv:2205.14949*, 2022. [2](#)
- [82] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. iBOT: Image bert pre-training with online tokenizer. In *ICLR*, 2022. [2](#), [7](#)