

Reduce Information Loss in Transformers for Pluralistic Image Inpainting

Qiankun Liu^{1*} Zhentao Tan¹ Dongdong Chen² Qi Chu^{1†} Xiyang Dai²

Yinpeng Chen² Mengchen Liu² Lu Yuan² Nenghai Yu¹

¹CAS Key Laboratory of Electromagnetic Space Information,
 University of Science and Technology of China

²Microsoft Cloud + AI

{liuqk3, tzt}@mail.ustc.edu.cn, {qchu, ynh}@ustc.edu.cn

cddlyf@gmail.com, {xiyang.dai, yiche, mengchliu, luyuan}@microsoft.com

Abstract

Transformers have achieved great success in pluralistic image inpainting recently. However, we find existing transformer based solutions regard each pixel as a token, thus suffer from information loss issue from two aspects: 1) They downsample the input image into much lower resolutions for efficiency consideration, incurring information loss and extra misalignment for the boundaries of masked regions. 2) They quantize 256^3 RGB pixels to a small number (such as 512) of quantized pixels. The indices of quantized pixels are used as tokens for the inputs and prediction targets of transformer. Although an extra CNN network is used to upsample and refine the low-resolution results, it is difficult to retrieve the lost information back. To keep input information as much as possible, we propose a new transformer based framework “PUT”. Specifically, to avoid input downsampling while maintaining the computation efficiency, we design a patch-based auto-encoder P-VQVAE, where the encoder converts the masked image into non-overlapped patch tokens and the decoder recovers the masked regions from the inpainted tokens while keeping the unmasked regions unchanged. To eliminate the information loss caused by quantization, an Un-Quantized Transformer (UQ-Transformer) is applied, which directly takes the features from P-VQVAE encoder as input without quantization and regards the quantized tokens only as prediction targets. Extensive experiments show that PUT greatly outperforms state-of-the-art methods on image fidelity, especially for large masked regions and complex large-scale datasets.

1. Introduction

Image inpainting, which focuses on filling meaningful and plausible contents in missing regions for the damaged

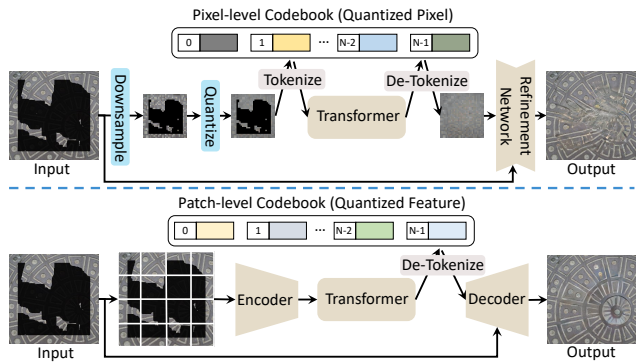


Figure 1. Top: Existing transformer based methods [49]. The output is produced by ICT [49]. Bottom: Our transformer based method. “Tokenize” here means getting the indices of quantized pixels or features, and “De-Tokenize” is the inverse operation.

images, has always been a hot topic in computer vision areas and widely used in various applications [1, 35, 39, 42, 47, 48, 56]. Traditional methods [1, 3, 9] based on texture matching can handle simple cases very well but struggle for complex natural images. In the last several years, benefiting from development of CNNs, tremendous success [28, 30, 34, 54] has been achieved by learning on large-scale datasets. However, due to the inherent properties of CNNs, i.e., local inductive bias and spatial-invariant kernels, such methods still do not perform well in understanding global structure and inpainting large masked/missing regions.

Recently, transformers have demonstrated their power in various vision tasks [4–8, 13, 16, 37, 50], thanks to their capability of modeling long-term relationship. Some recent works [49] also attempt to apply transformers for pluralistic image inpainting, and have achieved remarkable success in better diversity and large region inpainting quality. As shown in the top row of Figure 1, they follow the similar design: 1) Downsample the input image into lower resolutions and quantize the pixels; 2) Use the transformer to re-

*Work done during an internship at Microsoft

†Corresponding author

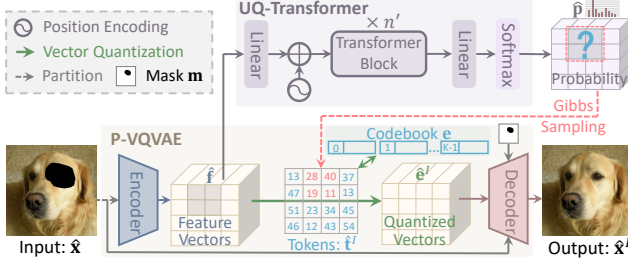


Figure 2. Pipeline of PUT for pluralistic image inpainting. Note that the pluralistic inpainting results are not shown in the figure.

cover the masked pixels by regarding each quantized pixel as the token; 3) Upsample and refine the low-resolution result by feeding it together with the original input image into an extra CNN network.

In this paper, we argue that using the above pixel-based token makes existing transformer based solutions suffer from the information loss issue from two aspects: 1) “*Low resolution*”. To avoid high computation complexity of transformer, the input image is downsampled into much lower resolution to reduce the input token number, which not only incurs information loss but also introduces misalignment for the boundaries of masked regions when upsampled back to the original resolution. 2) “*Quantization*”. To constrain the prediction within a small space, the huge amount (256^3 , in detail) of RGB pixels are quantized into much less (such as 512) quantized pixels through clustering. The indices of quantized pixels are used as discrete tokens both for the input and prediction target of transformer. Such practice would further result in the information loss.

To mitigate these issues, we propose a new transformer based framework **PUT**, which can reduce the information loss as much as possible. As shown in the bottom row of Figure 1, the original high-resolution input image is directly fed into a patch based encoder without any downsampling and the transformer directly takes the features from the encoder as input without any quantization. Specifically, **PUT** contains two key designs: Patch-based Vector Quantized Variational Auto-Encoder (“*P-VQVAE*”, Section 3.1) and Un-Quantized Transformer (“*UQ-Transformer*”, Section 3.2). *P-VQVAE* is a specially designed patch auto-encoder: 1) Its encoder converts each image patch into the latent feature in a non-overlapped way, where the non-overlapped design is to avoid the disturbance between masked regions and unmasked regions; 2) As the prediction space of *UQ-Transformer*, a dual-codebook is built for patch feature tokenization, where masked patches and unmasked patches are separately represented by different codebooks; 3) The decoder in *P-VQVAE* not only recovers the masked image regions from the inpainted tokens but also maintains unmasked regions unchanged. For *UQ-Transformer*, it utilizes the quantized tokens of unmasked

patches as the prediction targets for masked patches, but takes the un-quantized feature vectors from the encoder as input. Compared to taking the quantized tokens as input, this design can avoid the information loss and help *UQ-Transformer* make more accurate predictions.

To demonstrate the superiority, we conduct extensive experiments on FFHQ [25], Places2 [61] and ImageNet [11]. The results show that our method outperforms CNN based pluralistic inpainting methods by a large margin on different evaluation metrics. Benefiting from less information loss, our method also achieves much higher fidelity than existing transformer based solutions, especially for large region inpainting and complex large-scale datasets.

2. Related Work

Auto-Encoders. Auto-encoders [21] is a kind of artificial neural network in semi-supervised and unsupervised learning. Among its subclasses, variational auto-encoders (VAE) [12, 27] is widely used for image synthesis tasks [41, 43] as a generative model. It can be trained in self-supervised strategy and generate diverse images through decoder with latent space sampling or autoregressive models [33, 45]. Later, vector quantized variational auto-encoder (VQ-VAE) [44] is proposed for discrete representation learning to circumvent issues of “posterior collapse”, and further developed by VQ-VAE-2 [38]. Recently, based on the similar quantization mechanism with VQ-VAE, VAGAN [16] and dVAE [37] are proposed for conditional image generation through transformers [46], while PeCo [14] trains a perceptual vision tokenizer for vision transformer BERT-pretraining [51]. Different from previous methods, the proposed “*P-VQVAE*”, which contains a non-overlapped patch encoder, a dual-codebook and a multi-scale guided decoder, is dedicated for image inpainting.

Visual Transformers. Thanks to the capability of long range relationship modeling, transformers have been widely used in different vision tasks, such as object detection [4, 6], image synthesis [5, 16, 37], object tracking [7, 50] and image inpainting [49]. Specifically, the autoregressive inference mechanism is naturally suitable for image synthesis related tasks, which can bring diverse results while guarantee the quality of the synthesized images [5, 16, 37, 49]. In this paper, we take full advantage of transformer and propose to replace discrete tokens with continuous feature vectors to avoid the information loss.

Image Inpainting. According to the diversity of inpainted images, there are two different types of definition for image inpainting task: deterministic image inpainting and pluralistic image inpainting. Most of the traditional methods, whether diffusion-based methods [2, 15] or path-based methods [1, 10, 19], can only generate single result for each input and may be failed while meeting large-area of miss-

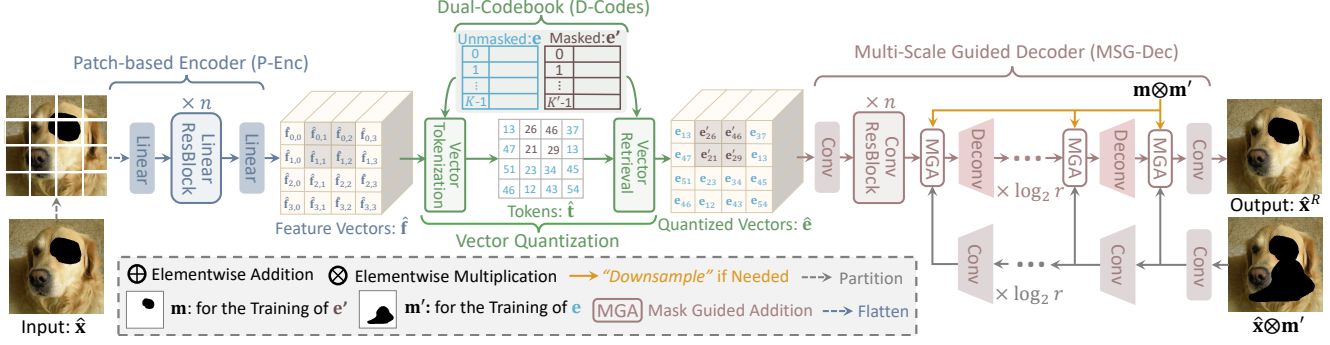


Figure 3. Training procedure of P-VQVAE. The detailed architecture of P-VQVAE could be found in the supplementary material.

ing pixels. Later, some CNN based methods [23, 28, 30, 32, 34, 54] are proposed to ensure consistency of the semantic content of the inpainted images, but still ignore the diversity of results. To generate several diverse results for each masked image, some CNN based [58, 60] and transformer based [49] methods have emerged recently. Among them, transformer based methods [49] show their merits in both quality and diversity than CNN based methods. However, their unreasonable design, like downsampling of input image and quantization of transformer inputs, results in the serious information loss issue. Thus, we propose a novel framework, PUT, which maximizes the input information to achieve better synthesis results.

3. Method

The proposed method mainly consists of a Patch-based Vector Quantized Variational Auto-Encoder (P-VQVAE) and an Un-Quantized transformer (UQ-Transformer). The overview of our method is shown in Figure 2. Let $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ be an image and $\mathbf{m} \in \{0, 1\}^{H \times W \times 1}$ be the mask denoting whether a region needs to be inpainted (with value 0) or not (with value 1). H and W are the spatial resolution. The image $\hat{\mathbf{x}} = \mathbf{x} \otimes \mathbf{m}$ is the masked image that contains missing pixels, where \otimes is the elementwise multiplication. The masked image $\hat{\mathbf{x}}$ is first fed into the encoder of P-VQVAE to get the patch based feature vectors. Then UQ-Transformer takes the feature vectors as input and predicts the tokens (i.e., indices) of latent vectors in a codebook for masked regions. Finally, the retrieved latent vectors are used as the quantized vectors for patches and fed into the decoder of P-VQVAE to reconstruct the inpainted image.

3.1. P-VQVAE

To avoid the information loss from input downsampling while maintaining the computation efficiency for the transformer, we utilize the merits of auto-encoder to replace the downsampled pixels with the features from encoder. Compared to the downsampled pixels, the features from encoder can have the same low-resolution for efficiency while con-

tains more information for reconstruction. Considering the task of image inpainting, we specially design P-VQVAE, which contains a patch-based encoder, a dual-codebook and a multi-scale guided decoder.

Patch-based Encoder. Conventional CNN based encoders process the input image with several convolution kernels in a sliding window manner, which are unsuitable for image inpainting since they would introduce disturbance between masked and unmasked regions. Thus, the encoder of P-VQVAE (denoted as P-Enc) is designed to process input image by several linear layers in a non-overlapped patch manner. Specifically, the masked image $\hat{\mathbf{x}}$ is firstly partitioned into $\frac{H}{r} \times \frac{W}{r}$ non-overlapped patches, where r is the spatial size of patches and is set to 8 by default. For a patch, we call it a *masked patch* if it contains any missing pixels, otherwise *unmasked patch*. Each patch is flattened and then mapped into a feature vector. Formally, all feature vectors are denoted as $\hat{\mathbf{f}} = \mathcal{E}(\hat{\mathbf{x}}) \in \mathbb{R}^{\frac{H}{r} \times \frac{W}{r} \times C}$, where C (set to 256 by default) is the dimensionality of feature vectors and $\mathcal{E}(\cdot)$ is the encoder function.

Dual-Codebook for Vector Quantization. Following the works in [16, 38, 44], the feature vectors from encoder are quantized into discrete tokens with the latent vectors in the learnable codebook. By contrast, we design a dual-codebook (denote as D-Codes) for vector quantization, which is more suitable for image inpainting. In D-Codes, the latent vectors are divided into two parts, denoted as $\mathbf{e} \in \mathbb{R}^{K \times C}$ and $\mathbf{e}' \in \mathbb{R}^{K' \times C}$, which are responsible for feature vectors that mapped from unmasked and masked patches respectively. K and K' are the number of latent vectors. Let $\mathbf{m}^\downarrow \in \{0, 1\}^{\frac{H}{r} \times \frac{W}{r} \times 1}$ be the indicator mask that indicates whether a patch is a masked (with value 0) or unmasked (with value 1) patch. The feature vector $\hat{\mathbf{f}}_{i,j}$ is quantized as:

$$\begin{cases} \mathbf{e}_k & \text{where } k = \operatorname{argmin}_l \|\hat{\mathbf{f}}_{i,j} \ominus \mathbf{e}_l\|_2, \text{ if } \mathbf{m}_{i,j}^\downarrow = 1, \\ \mathbf{e}'_{k'} & \text{where } k' = \operatorname{argmin}_l \|\hat{\mathbf{f}}_{i,j} \ominus \mathbf{e}'_l\|_2, \text{ else,} \end{cases} \quad (1)$$

where \ominus denotes the operation of elementwise subtraction. Let $\hat{\mathbf{e}} \in \mathbb{R}^{\frac{H}{r} \times \frac{W}{r} \times C}$ and $\hat{\mathbf{t}} = \mathcal{I}(\hat{\mathbf{f}}, \mathbf{e}, \mathbf{e}', \mathbf{m}^\downarrow) \in \mathbb{N}^{\frac{H}{r} \times \frac{W}{r}}$ be the quantized vectors and tokens for $\hat{\mathbf{f}}$, where $\mathcal{I}(\cdot, \cdot, \cdot, \cdot)$ denotes the function that gets tokens for its first argument and it can be simply implemented by getting the indices of all quantized vectors in $\hat{\mathbf{e}}$. The dual-codebook helps P-Enc learn more discriminative feature vectors for masked and unmasked patches since they are quantized and represented with different codebooks, which further disenchant transformer about the masked and unmasked patches to predict more reasonable results for masked patches.

Multi-Scale Guided Decoder. For image inpainting task, an indisputable fact is that the unmasked regions should be kept unchanged. To this end, we design a multi-scale guided decoder, denoted as MSG-Dec, to construct the inpainted image $\hat{\mathbf{x}}^I$ by referencing the input masked image $\hat{\mathbf{x}}$. Let $\hat{\mathbf{t}}^I$ be the inpainted tokens by transformer (Ref. Figure 2 and Section 3.3) and $\hat{\mathbf{e}}^I$ be the retrieved quantized vectors from codebook based on $\hat{\mathbf{t}}^I$. The construction procedure is formulated as:

$$\hat{\mathbf{x}}^I = \mathcal{D}(\hat{\mathbf{e}}^I, \mathbf{m}, \hat{\mathbf{x}}), \quad (2)$$

where $\mathcal{D}(\cdot, \cdot, \cdot)$ is the decoder function. The decoder consists of two branches: a main branch which starts with the quantized vectors $\hat{\mathbf{e}}^I$ and uses several deconvolution layers to generate inpainted images and a reference branch which extracts multi-scale feature maps (with spatial sizes $\frac{H}{2^l} \times \frac{W}{2^l}$, $0 \leq l \leq \log_2 r$) from the masked image $\hat{\mathbf{x}}$. The features from the reference branch are fused to the features with the same scale from the main branch through a Mask Guided Addition (MGA) module as:

$$\hat{\mathbf{e}}^{I,l-1} = \text{Deconv}((1 - \mathbf{m}^{\downarrow,l}) \otimes \hat{\mathbf{e}}^{I,l} + \mathbf{m}^{\downarrow,l} \otimes \hat{\mathbf{f}}^{R,l}), \quad (3)$$

where $\hat{\mathbf{e}}^{I,l}$ and $\hat{\mathbf{f}}^{R,l}$ are the features with spatial size $\frac{H}{2^l} \times \frac{W}{2^l}$ from the main branch and the reference branch respectively. $\mathbf{m}^{\downarrow,l}$ is the indicator mask obtained from \mathbf{m} for corresponding spatial size.

Training of P-VQVAE. To avoid the decoder learning to reconstruct input image $\hat{\mathbf{x}}$ only from the reference image, we get the reference image by randomly erasing some pixels in $\hat{\mathbf{x}}$ with another mask \mathbf{m}' (see in Figure 3). Let $\hat{\mathbf{x}}^R = \mathcal{D}(\hat{\mathbf{e}}, \mathbf{m} \otimes \mathbf{m}', \hat{\mathbf{x}} \otimes \mathbf{m}')$ be the reconstructed image. In our design, the unmasked pixels in the reference image will be utilized to recover the corresponding pixels in $\hat{\mathbf{x}}^R$, while the latent vectors in codebook \mathbf{e}' and \mathbf{e} will be used to recover the pixels in $\hat{\mathbf{x}}^R$ masked by \mathbf{m} and the remaining pixels respectively. The loss for training P-VQVAE is:

$$L_{vae} = \mathcal{L}_{rec}(\hat{\mathbf{x}}, \hat{\mathbf{x}}^R) + \|\text{sg}[\hat{\mathbf{f}}] \ominus \hat{\mathbf{e}}\|_2^2 + \beta \|\text{sg}[\hat{\mathbf{e}}] \ominus \hat{\mathbf{f}}\|_2^2, \quad (4)$$

where $\text{sg}[\cdot]$ refers to a stop-gradient operation that blocks gradients from flowing into its argument. β is the weight for balance and is set to 0.25. $\mathcal{L}_{rec}(\cdot, \cdot)$ is the function to measure the difference between inputted and reconstructed images, including the L1 loss between the pixel values in two images and the gradients of two images, the adversarial loss [18] obtained by a discriminator network, as well as the perceptual loss [24] and the style loss [17] between the two images. Following [38, 44], the second term in Eq. (4) is replaced by Exponential Moving Average (EMA) to optimize the vectors in D-Codes. More details about the training of P-VQVAE could be found in the supplementary material.

3.2. UQ-Transformer

In existing transformers for image inpainting [49] and synthesis [16, 37], the quantized discrete tokens are used as both the inputs and prediction targets. Given such discrete tokens, transformers suffer from the severe information loss issue, which is harmful to their prediction. In contrast, to take full advantage of feature vectors $\hat{\mathbf{f}}$ from the encoder of P-VQVAE, our UQ-Transformer directly takes them as the inputs and predicts the discrete tokens for masked patches.

Specifically, $\hat{\mathbf{f}}$ is firstly mapped by a linear layer and then added with extra learnable position embeddings for the encoding of spatial information. Finally, following [36], the feature vectors are flattened along spatial dimension to get the final input for the subsequent several transformer blocks. The output of the last transformer block is further projected to the distribution over K latent vectors in codebook \mathbf{e} with a linear layer and a softmax function. We formulate the above procedure as $\hat{\mathbf{p}} = \mathcal{T}(\hat{\mathbf{f}}) \in [0, 1]^{\frac{H}{r} \times \frac{W}{r} \times K}$, where $\mathcal{T}(\cdot)$ refers the UQ-Transformer function.

Training of UQ-Transformer. Given a masked image $\hat{\mathbf{x}}$, the distribution of its corresponding inpainted tokens over K latent vectors can be obtained with the pre-trained P-VQVAE and UQ-Transformer $\hat{\mathbf{p}} = \mathcal{T}(\mathcal{E}(\hat{\mathbf{x}}))$. The ground-truth tokens for \mathbf{x} is $\mathbf{t} = \mathcal{I}(\mathcal{E}(\mathbf{x}), \mathbf{e}, \mathbf{e}', \mathcal{O}(\mathbf{m}^\downarrow))$ (Ref. Section 3.1), where $\mathcal{O}(\cdot)$ sets all values in the given argument to 1. UQ-Transformer is trained with cross-entropy loss by fixing P-VQVAE:

$$L_{trans} = \frac{-1}{\sum_{i,j} 1 - \mathbf{m}_{i,j}^\downarrow} \sum_{i,j} (1 - \mathbf{m}_{i,j}^\downarrow) \log \hat{\mathbf{p}}_{i,j, \mathbf{t}_{i,j}}. \quad (5)$$

In order to make the training stage consistent with inference stage, where only the quantized vectors can be obtained for masked regions, we randomly quantize the feature vectors in $\mathcal{E}(\hat{\mathbf{x}})$ to the latent vectors in codebook with probability 0.3 before feeding it to UQ-Transformer.

3.3. Sampling Strategy for Image Inpainting

For the production of diverse results, the tokens for masked patches ($\mathbf{m}_{i,j}^\downarrow = 0$) are iteratively sampled with

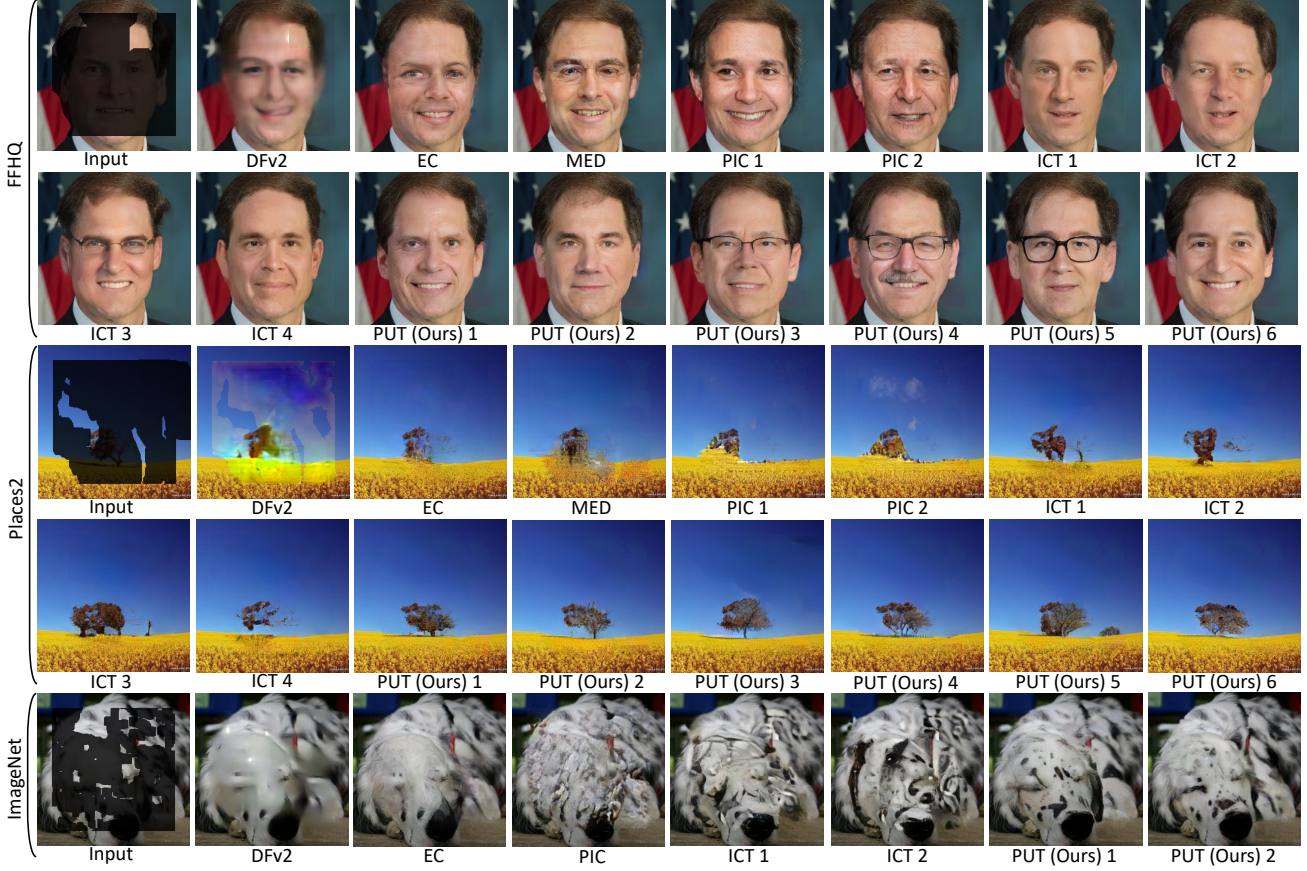


Figure 4. Samples of inpainted results produced by different methods. For PUT, we set $\mathcal{K} = 50$. More qualitative comparisons are presented in the supplementary material.

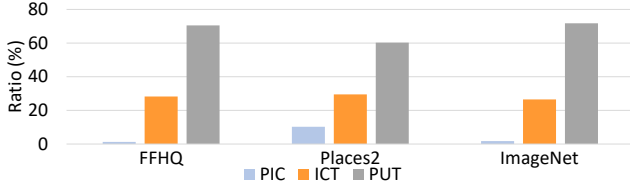


Figure 5. The ratio of each method among the rank 1 images evaluated by human. Statistics are collected from 23 participants.

Gibbs sampling. Specifically, in each iteration, we first select the patch with the maximum predicted probability among the remaining masked patches. Then the token for the selected patch is sampled from the top- \mathcal{K} predicted elements. Finally, the corresponding latent vector for the sampled token is retrieved to replace the feature vector of the selected patch before feeding UQ-Transformer for the next iteration. After sampling the tokens for all masked patches, we can get all quantized vectors $\hat{\mathbf{e}}^I$ with the inpainted tokens $\hat{\mathbf{t}}^I$, and the inpainted image can be constructed using Eq. (2). For the production of deterministic results, the tokens for masked patches are sampled with the largest probabilities at once.

4. Experiments

The evaluation is conducted at 256×256 (i.e., $H = 256$ and $W = 256$) resolution on three different datasets, including FFHQ [25], Places2 [61] and ImageNet [11]. We use the original training and testing splits for Places2 and ImageNet. For FFHQ, we maintain the last 1K images for evaluation and others for training. Following ICT [49], only 1K images are randomly chosen from the test split of ImageNet for evaluation and the irregular masks provided by PConv [28] are used both for training and testing.

4.1. Implementation Details

We use P-VQVAE with the same model size and UQ-Transformer with different model sizes for different datasets. The number of latent vectors in dual-codebook (i.e., K and K') both are set to 512. The detailed architecture of P-VQVAE and UQ-Transformer could be found in the supplementary material. We train P-VQVAE with batch size 128 and train UQ-Transformer with batch size 48 for FFHQ and 96 for Places2 and ImageNet. The learning rate is warmed up from 0 to $2e-4$ and $3e-4$ in the first 5K iterations for P-VQVAE and UQ-Transformer, and then de-

Dataset		FFHQ [25]			Places2 [61]			ImageNet [11]		
Mask Ratio (%)		20-40	40-60	10-60	20-40	40-60	10-60	20-40	40-60	10-60
FID ↓	DFv2 (ICCV, 2019) [55]	27.344	47.894	30.509	53.107	83.979	59.280	49.900	102.111	64.056
	EC (ICCVW, 2019) [32]	12.949	26.217	16.961	20.180	34.965	23.206	27.821	63.768	39.199
	MED (ECCV, 2020) [29]	13.999	26.252	17.061	28.671	46.815	32.494	40.643	93.983	54.854
	ICT _{all} (ICCV, 2021) [49]	10.442	23.946	15.363	19.309	33.510	23.331	23.889	54.327	32.624
	PUT _{all} (Ours)	11.221	19.934	13.248	19.776	38.206	24.605	19.411	43.239	26.223
	PIC (CVPR, 2019) [60]	22.847	37.762	25.902	31.361	44.289	34.520	49.215	102.561	63.955
	ICT ₅₀ (ICCV, 2021) [49]	13.536	23.756	16.202	20.900	33.696	24.138	25.235	55.598	34.247
PSNR ↑	PUT ₅₀ (Ours)	12.784	21.382	14.554	19.617	31.485	22.121	21.272	45.153	27.648
	DFv2 (ICCV, 2019) [55]	27.937	22.984	26.783	26.292	22.412	25.391	24.464	20.157	23.387
	EC (ICCVW, 2019) [32]	27.484	22.574	26.181	26.536	22.755	25.975	24.703	20.459	23.596
	MED (ECCV, 2020) [29]	27.117	22.499	26.111	25.401	21.543	24.510	23.730	19.560	22.752
	ICT _{all} (ICCV, 2021) [49]	29.847	23.041	26.736	25.836	22.120	24.986	24.249	20.045	23.317
	PUT _{all} (Ours)	28.356	24.125	27.473	26.580	22.945	25.749	25.721	21.551	24.726
	PIC (CVPR, 2019) [60]	25.157	20.424	24.093	24.073	20.656	23.469	22.921	18.368	21.623
SSIM ↑	ICT ₅₀ (ICCV, 2021) [49]	26.462	21.816	25.515	24.947	21.126	24.373	23.252	19.025	22.123
	PUT ₅₀ (Ours)	26.877	22.375	25.943	25.452	21.528	24.492	24.238	19.742	23.264
	DFv2 (ICCV, 2019) [55]	0.945	0.850	0.912	0.878	0.741	0.831	0.876	0.719	0.819
	EC (ICCVW, 2019) [32]	0.941	0.826	0.899	0.881	0.734	0.840	0.882	0.714	0.824
	MED (ECCV, 2020) [29]	0.936	0.840	0.903	0.854	0.685	0.796	0.861	0.675	0.795
	ICT _{all} (ICCV, 2021) [49]	0.964	0.863	0.917	0.870	0.723	0.819	0.876	0.711	0.818
	PUT _{all} (Ours)	0.953	0.888	0.908	0.885	0.756	0.840	0.904	0.772	0.838
MAE ↓	PIC (CVPR, 2019) [60]	0.910	0.769	0.865	0.824	0.648	0.775	0.842	0.623	0.766
	ICT ₅₀ (ICCV, 2021) [49]	0.931	0.822	0.896	0.850	0.682	0.803	0.852	0.666	0.786
	PUT ₅₀ (Ours)	0.936	0.845	0.906	0.861	0.703	0.806	0.875	0.704	0.818
	DFv2 (ICCV, 2019) [55]	0.0187	0.0429	0.0270	0.0230	0.0461	0.0304	0.0303	0.0638	0.0415
	EC (ICCVW, 2019) [32]	0.0177	0.0430	0.0263	0.0207	0.0419	0.0261	0.0271	0.0582	0.0375
	MED (ECCV, 2020) [29]	0.0200	0.0430	0.0277	0.0255	0.0505	0.0336	0.0320	0.0676	0.0434
	ICT _{all} (ICCV, 2021) [49]	0.0129	0.0368	0.0232	0.0221	0.0433	0.0289	0.0362	0.0578	0.0378
MAE ↓	PUT _{all} (Ours)	0.0159	0.0328	0.0213	0.0205	0.0398	0.0269	0.0233	0.0487	0.0321
	PIC (CVPR, 2019) [60]	0.0251	0.0571	0.0350	0.0284	0.0544	0.0353	0.0361	0.0785	0.0509
	ICT ₅₀ (ICCV, 2021) [49]	0.0196	0.0445	0.0270	0.0245	0.0487	0.0312	0.0312	0.0677	0.0440
MAE ↓	PUT ₅₀ (Ours)	0.0191	0.0417	0.0263	0.0235	0.0479	0.0317	0.0281	0.0641	0.0401

Table 1. Quantitative results of different methods. The methods are divided into deterministic and pluralistic groups. The subscript “50” of ICT and PUT is the value of \mathcal{K} , while the subscript “all” of ICT and PUT means all tokens are sampled at one iteration.

cayed with cosine scheduler. P-VQVAE is optimized with Adam [26] ($\beta_1 = 0, \beta_2 = 0.9$) and UQ-Transformer is optimized with AdamW [31] ($\beta_1 = 0.9, \beta_2 = 0.95$). All models are trained to their convergence.

4.2. Main Results

We compare the proposed PUT with the following state-of-the-art inpainting approaches: DeepFillv2 (DFv2) [55], Edge-Connect (EC) [32], MED [29], PIC [60] and ICT [49]. Among them, the last two ones can generate pluralistic results for each input while the others can only produce one deterministic result for the given input. For a fair comparison, we directly use the pre-trained models provided by the authors when available, otherwise train the models by ourselves using the codes and settings provided by the authors.

Qualitative Comparisons. We first qualitatively compare the inpainted results of different methods in Figure 4. Specifically, DFv2 and EC generally produce blurry images, while the texture of the results generated by MED and PIC contain lots of unnatural artifacts. Compared with ICT, PUT is more powerful in the understanding of global context and maintaining the meaningful textures, as shown

in the last row in Figure 4. We argue that the superiority of PUT is mainly due to: 1) high-resolution and 2) unquantized transformer. Both of them are vital for preserving the information contained in the input images, which are helpful to produce photo-realism images.

User Study. For the evaluation of subjective quality, the user study is further conducted. Only pluralistic methods are evaluated, including PIC [60], ICT [49] and PUT. Specifically, we randomly sample 20 pairs of image and mask from the test set of each dataset. For each pair, we generate two inpainted results using each method, and ask the participants to rank these six images according to their photo-realism from high to low. We calculate the ratio of each method among the rank 1 images. Results are shown in Figure 5. Our method takes up at least 60% of the rank 1 images, demonstrating its superiority.

Quantitative Comparisons. We further demonstrate the superiority of PUT in diversity and fidelity with other pluralistic methods. Specifically, the mean LPIPS distance [57] between pairs of randomly generated results for the same input image is calculated. Following ICT [49], five pairs

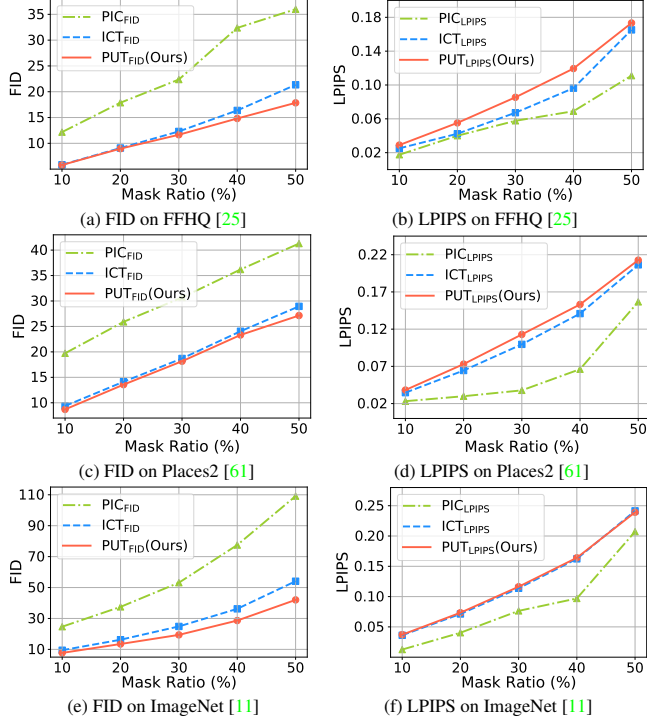


Figure 6. LPIPS and FID curves with respect to mask ratio on different datasets. For PUT and ICT, $\mathcal{K} = 50$.

per input image are generated. Meanwhile, the Fréchet Inception Distance (FID) [20] is also computed between the inpainted images and ground-truth images to reveal the fidelity of generated results. The curves of LPIPS and FID are shown in Figure 6. It can be seen that among the three pluralistic methods, PUT achieves the best fidelity (lowest FID) on all datasets, especially for large area of masked regions and complex scenes (ImageNet [11]). Although PUT and ICT both are implemented with transformers, PUT presents a higher diversity, we owe this to the feeding of original continuous feature vectors without quantization to transformer, which has no information loss issue.

In Table 1, we compare all methods in terms of several metrics, including the peak signal-to-noise (PSNR), structural similarity index (SSIM), relative L1 (MAE) and FID. Only one recovered output is produced for each input. Overall, for pluralistic methods, PUT₅₀ performs the best in almost all metrics on all datasets. Specifically, the FID score of PUT₅₀ on ImageNet with mask ratio 40%-60% is 10.44 lower than that of ICT₅₀. For deterministic methods, PUT_{all} also achieves almost the best performance. However, the metrics are suitable for the comparison between pluralistic and deterministic methods since diverse and meaningful results can be generated by pluralistic methods.

In Table 2, we further compare PUT_{all} with LaMa [40] and CoModGAN [59], both of which are recently proposed methods for image inpainting with resolution 512×512 . For

Dataset	Metric	PUT _{all} (Ours)	CoModGAN [59]	LaMa [40]
FFHQ [25]	PSNR \uparrow / FID \downarrow	24.245/21.351	22.430/ 17.914	-/-
Places2 [61]	PSNR \uparrow / FID \downarrow	22.589/38.472	20.962/32.559	22.694/31.160

Table 2. Performance of different methods on FFHQ [25] and Places2 [61] with mask ratio 40%-60% and resolution 512×512 .

Dataset		FFHQ [25]			Places2 [61]		
Mask Ratio (%)		20-40	40-60	10-60	20-40	40-60	10-60
FID \downarrow	PUT ^{conv}	163.610	226.437	173.351	178.057	216.235	179.294
	PUT ^{one}	12.112	20.298	13.960	26.015	38.011	28.634
	PUT ^{no.ref}	15.014	22.736	16.469	22.821	32.281	25.084
	PUT ^{tok}	22.824	38.384	26.098	58.643	120.482	75.625
	PUT ^{qua0}	43.621	96.648	54.879	35.127	71.629	44.588
	PUT	12.784	21.382	14.554	19.617	31.485	22.121
PSNR \uparrow	PUT ^{conv}	12.783	9.735	12.360	12.207	9.347	11.799
	PUT ^{one}	26.887	22.335	25.903	24.507	20.571	23.600
	PUT ^{no.ref}	26.487	22.249	25.547	25.053	21.398	24.185
	PUT ^{tok}	23.916	18.936	22.879	20.940	14.685	19.429
	PUT ^{qua0}	24.188	19.564	23.174	24.340	20.239	23.353
	PUT	26.877	22.375	25.943	25.452	21.528	24.492
SSIM \uparrow	PUT ^{conv}	0.495	0.266	0.445	0.417	0.212	0.373
	PUT ^{one}	0.937	0.844	0.906	0.836	0.660	0.776
	PUT ^{no.ref}	0.932	0.841	0.901	0.851	0.695	0.798
	PUT ^{tok}	0.893	0.745	0.843	0.742	0.453	0.652
	PUT ^{qua0}	0.879	0.704	0.820	0.831	0.637	0.763
	PUT	0.945	0.857	0.914	0.861	0.703	0.806
MAE \downarrow	PUT ^{conv}	0.1420	0.3017	0.1912	0.1447	0.2995	0.1941
	PUT ^{one}	0.0191	0.0418	0.0264	0.0256	0.0526	0.0345
	PUT ^{no.ref}	0.0202	0.0426	0.0274	0.0249	0.0488	0.0329
	PUT ^{tok}	0.0289	0.0689	0.0418	0.0467	0.1337	0.0779
	PUT ^{qua0}	0.0281	0.0627	0.0393	0.0269	0.0562	0.0367
	PUT	0.0191	0.0417	0.0263	0.0235	0.0479	0.0317

Table 3. Quantitative results of different methods. All the results are tested with $\mathcal{K} = 50$.

PUT_{all}, an upsample network used in ICT [49] is trained to upsample the results from 256×256 to 512×512 . As we can see, PUT_{all} achieves a much better PSNR than CoModGAN, and it is also comparable with LaMa.

4.3. Discussions

Effectiveness of different components. To show the effectiveness of the patch-based encoder, dual-codebook, multi-scale guided decoder and un-quantized transformer, several methods are designed: 1) PUT^{conv} means that the patch-based encoder is replaced with a normal CNN encoder, which is implemented with convolution layers; 2) PUT^{one} refers that there is only one codebook for vector quantization; 3) PUT^{no.ref} means that there is no reference branch in the decoder; and 4) PUT^{tok} denotes that the transformer takes the quantized vectors as input rather than the original feature vectors from encoder. In order to show the effectiveness of random quantization of feature vectors while training UQ-Transformer, we further design the fifth model, PUT^{qua0}, that trains UQ-Transformer without random quantization. For all models, except the modifications mentioned, others remain the same with our default model




Figure 7. Inpainted results of different models on Places2 [61].

PUT. Results are shown in Table 3 and Figure 7.

Among all models, PUT^{conv} performs the worst in all metrics, demonstrating the effectiveness of the non-overlapping patch partition design. Within CNN based encoder, the input images are processed in a sliding window manner, introducing the interaction between masked and unmasked regions, which is fatal to transformer for the prediction of masked regions.

Comparing PUT^{one} to PUT, the only difference is the training of P-VQVAE with one or two codebooks since the codebook e' is not used in the inference stage. P-VQVAE can learn more discriminative features for masked and unmasked patches with the help of dual-codebook. Interestingly, PUT indeed performs better than PUT^{one} except the FID score on FFHQ. We speculate that face generation is much easier because all faces share a similar structure. Besides, the facial structure is related with the position in the image since most of the images in FFHQ contain faces in their near-center locations. As we can see in Figure 7, PUT^{one} sometimes predicts *black* patches, which is similar to those patches containing missing pixels. Nonetheless, PUT achieves overall better performance than PUT^{one} .

Compared with PUT, PUT^{no-ref} constructs the inpainted image without referencing to the input masked image, which leads to a inferior performance in terms of all metrics. For PUT, some useful textures can be recovered with the help of the guidance from the unmasked regions in reference image. As we can see in Figure 7, the result of PUT^{no-ref} is over smoothed, which is unnatural.

Compared with PUT^{tok} , PUT performs much better in all metrics. Without quantizing feature vectors to discrete representations, no information contained in the feature will be lost. Such practice helps transformer to understand complex content and maintain the inherit meaningful textures in the input image. However, the training of UQ-Transformer should be carefully designed by randomly quantizing the input feature vectors since only quantized vectors can be obtained for masked regions at the inference stage. The effectiveness of such random quantization during training is obvious while comparing PUT with PUT^{qua0} .




Figure 8. Misalignment between the high and low resolutions. The mask and input image in low-resolution are zoomed in for a better comparison.

Misalignment between high and low resolutions. Existing transformer based methods [49] involves high and low resolutions. The images and masks in high and low resolutions are misaligned, especially the borders of masked regions. For real applications, where the masked regions are often with arbitrary shapes and sizes, such misalignment is non-negligible. In Figure 8, we show one case that with irregular masked regions. We can see that the results of ICT [49] contain lots of artifacts near the misaligned borders. For PUT, the generated results are natural and smooth even though the provided mask is irregular.

5. Conclusions and Limitations

In this paper, we present a novel method, PUT, for pluralistic image inpainting. PUT consists of two main components: 1) patch-based auto-encoder (P-VQVAE) and 2) unquantized transformer (UQ-Transformer). With the help of P-VQVAE and UQ-Transformer, PUT processes the original high-resolution image without quantization. Such practice preserves the information contained in the input image as much as possible. The main limitation of PUT is the inference speed for the production of diverse results. However, it is a common issue of existing transformer based auto-regressive methods [16, 37, 46, 49]. Two solutions could be adopted to alleviate this limitation: 1) replacing the used transformer block with more efficient ones [22, 52] and 2) sampling tokens for several patches at each iteration. In addition, PUT may be used for editing the contents of images to achieve illicit goals, which can be mitigated using existing synthesized image detectors [53]. Experimental results demonstrate the superiority of PUT, including the fidelity and diversity, especially for large masked regions and complex scenes (such as ImageNet [11]).

Acknowledgement. This work is supported by the National Natural Science Foundation of China (No. 62002336, No. U20B2047) and the Exploration Fund Project of University of Science and Technology of China (YD3480002001).

References

- [1] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009. 1, 2
- [2] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 417–424, 2000. 2
- [3] Marcelo Bertalmio, Luminita Vese, Guillermo Sapiro, and Stanley Osher. Simultaneous structure and texture image inpainting. *IEEE transactions on image processing*, 12(8):882–889, 2003. 1
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020. 1, 2
- [5] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *International Conference on Machine Learning*, pages 1691–1703. PMLR, 2020. 1, 2
- [6] Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852*, 2021. 1, 2
- [7] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8126–8135, 2021. 1, 2
- [8] Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Xiaoyi Dong, Lu Yuan, and Zicheng Liu. Mobileformer: Bridging mobilenet and transformer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2022)*, 2022. 1
- [9] Antonio Criminisi, Patrick Pérez, and Kentaro Toyama. Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on image processing*, 13(9):1200–1212, 2004. 1
- [10] Soheil Darabi, Eli Shechtman, Connelly Barnes, Dan B Goldman, and Pradeep Sen. Image melding: Combining inconsistent images using patch-based synthesis. *ACM Transactions on graphics (TOG)*, 31(4):1–10, 2012. 2
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2, 5, 6, 7, 8
- [12] Carl Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016. 2
- [13] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. *arXiv preprint arXiv:2107.00652*, 2021. 1
- [14] 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
- [15] Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 341–346, 2001. 2
- [16] 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. 1, 2, 3, 4, 8
- [17] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 4
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 4
- [19] James Hays and Alexei A Efros. Scene completion using millions of photographs. *ACM Transactions on Graphics (ToG)*, 26(3):4-es, 2007. 2
- [20] 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. 7
- [21] Geoffrey E Hinton and Richard S Zemel. Autoencoders, minimum description length, and helmholtz free energy. *Advances in neural information processing systems*, 6:3–10, 1994. 2
- [22] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019. 8
- [23] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):1–14, 2017. 3
- [24] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 4
- [25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 2, 5, 6, 7
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [27] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [28] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 85–100, 2018. 1, 3, 5

- [29] Hongyu Liu, Bin Jiang, Yibing Song, Wei Huang, and Chao Yang. Rethinking image inpainting via a mutual encoder-decoder with feature equalizations. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 725–741. Springer, 2020. 6
- [30] Hongyu Liu, Bin Jiang, Yi Xiao, and Chao Yang. Coherent semantic attention for image inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4170–4179, 2019. 1, 3
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [32] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. Edgeconnect: Structure guided image inpainting using edge prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3, 6
- [33] Aäron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional image generation with pixelcnn decoders. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 4797–4805, 2016. 2
- [34] 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. 1, 3
- [35] Haonan Qiu, Chaowei Xiao, Lei Yang, Xinchun Yan, Honglak Lee, and Bo Li. Semanticadv: Generating adversarial examples via attribute-conditioned image editing. In *European Conference on Computer Vision*, pages 19–37. Springer, 2020. 1
- [36] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 4
- [37] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021. 1, 2, 4, 8
- [38] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in neural information processing systems*, pages 14866–14876, 2019. 2, 3, 4
- [39] Yujing Sun, Yizhou Yu, and Wenping Wang. Moiré photo restoration using multiresolution convolutional neural networks. *IEEE Transactions on Image Processing*, 27(8):4160–4172, 2018. 1
- [40] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2149–2159, 2022. 7
- [41] Zhentao Tan, Menglei Chai, Dongdong Chen, Jing Liao, Qi Chu, Bin Liu, Gang Hua, and Nenghai Yu. Diverse semantic image synthesis via probability distribution modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7962–7971, 2021. 2
- [42] Zhentao Tan, Menglei Chai, Dongdong Chen, Jing Liao, Qi Chu, Lu Yuan, Sergey Tulyakov, and Nenghai Yu. Michigan: multi-input-conditioned hair image generation for portrait editing. *ACM Transactions on Graphics (TOG)*, 39(4):95–1, 2020. 1
- [43] Zhentao Tan, Dongdong Chen, Qi Chu, Menglei Chai, Jing Liao, Mingming He, Lu Yuan, Gang Hua, and Nenghai Yu. Efficient semantic image synthesis via class-adaptive normalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [44] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6309–6318, 2017. 2, 3, 4
- [45] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756. PMLR, 2016. 2
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2, 8
- [47] Ziyu Wan, Bo Zhang, Dongdong Chen, Pan Zhang, Dong Chen, Jing Liao, and Fang Wen. Bringing old photos back to life. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2747–2757, 2020. 1
- [48] Ziyu Wan, Bo Zhang, Dongdong Chen, Pan Zhang, Dong Chen, Jing Liao, and Fang Wen. Old photo restoration via deep latent space translation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1
- [49] Ziyu Wan, Jingbo Zhang, Dongdong Chen, and Jing Liao. High-fidelity pluralistic image completion with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4692–4701, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [50] Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1571–1580, 2021. 1, 2
- [51] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luowei Zhou, and Lu Yuan. Bevt: Bert pretraining of video transformers. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2022)*, 2022. 2
- [52] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. 8
- [53] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8695–8704, 2020. 8

- [54] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5505–5514, 2018. 1, 3
- [55] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4471–4480, 2019. 6
- [56] Xiaohang Zhan, Xingang Pan, Bo Dai, Ziwei Liu, Dahua Lin, and Chen Change Loy. Self-supervised scene de-occlusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3784–3792, 2020. 1
- [57] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [58] Lei Zhao, Qihang Mo, Sihuan Lin, Zhizhong Wang, Zhiwen Zuo, Haibo Chen, Wei Xing, and Dongming Lu. Uctgan: Diverse image inpainting based on unsupervised cross-space translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5741–5750, 2020. 3
- [59] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I Chang, and Yan Xu. Large scale image completion via co-modulated generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2021. 7
- [60] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. Pluralistic image completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1438–1447, 2019. 3, 6
- [61] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017. 2, 5, 6, 7, 8