

# Incremental Transformer Structure Enhanced Image Inpainting with Masking Positional Encoding

Qiaole Dong\*, Chenjie Cao\*, Yanwei Fu†  
 School of Data Science, Fudan University

{18307130096, 20110980001, yanweifu}@fudan.edu.cn

## Abstract

Image inpainting has made significant advances in recent years. However, it is still challenging to recover corrupted images with both vivid textures and reasonable structures. Some specific methods only tackle regular textures while losing holistic structures due to the limited receptive fields of convolutional neural networks (CNNs). On the other hand, attention-based models can learn better long-range dependency for the structure recovery, but they are limited by the heavy computation for inference with large image sizes. To address these issues, we propose to leverage an additional structure restorer to facilitate the image inpainting incrementally. The proposed model restores holistic image structures with a powerful attention-based transformer model in a fixed low-resolution sketch space. Such a grayscale space is easy to be upsampled to larger scales to convey correct structural information. Our structure restorer can be integrated with other pretrained inpainting models efficiently with the zero-initialized residual addition. Furthermore, a masking positional encoding strategy is utilized to improve the performance with large irregular masks. Extensive experiments on various datasets validate the efficacy of our model compared with other competitors. Our codes are released in [https://github.com/DQiaole/ZITS\\_inpainting](https://github.com/DQiaole/ZITS_inpainting).

## 1. Introduction

Image inpainting has been investigated as a long-standing challenge to address the difficulty of filling in missing areas of pictures. It is very useful to various real-world applications, such as object removal [12], photo restoration, and image editing [25]. To achieve realistic outcomes, the inpainted images should remain both semantically coherent textures and visually reasonable structures. Many classical algorithms [3, 9, 19, 31, 42] search similar patches for the reconstruction heuristically. But preserving

\*Equal contributions.

†Corresponding authors.



Figure 1. High quality 1024×1024 inpainted results. From left to right, masked inputs, results of LaMa [44], results of our method.

good textures and holistic structures in large images is still non-trivial for these conventional methods.

Benefited from excellent capacities of Convolutional Neural Networks (CNNs) [28] and Generative Adversarial Networks (GANs) [15], existing deep learning methods [4, 18, 30, 35, 44, 46, 54, 59] could efficiently conduct the image inpainting tasks in some common cases. However, they still suffer from some dilemmas. (1) *Limited receptive fields*. Learning semantically consistent textures is difficult for traditional CNNs due to the local inductive priors and narrow receptive field of convolution operations. Even dilated convolutions [55] fail to tackle large corrupted regions or high-resolution images. (2) *Missing holistic structures*. Recovering key edges and lines for scenes, especially ones with weak texture is difficult without the holistic understanding of large images as shown in Fig. 1. (3) *Heavy computations*. Training GANs with large image sizes is still very tricky and costly [27]. And the inpainting performance may be degraded on high-resolution images. (4) *No positional information in masked regions*. The inpainting model tends to repeat meaningless artifacts in large irregular masked regions without explicit positional clues.

Some pioneering works can partially solve these problems. For the limited receptive fields, attention-based meth-

ods [54, 56, 60] leverage the attention mechanism to extend the receptive fields. Suvorov *et al.* [44] utilize the Fast Fourier Convolution (FFC) to encode features in frequency fields with global receptive fields for resolution-robust inpainting. But they fail to ensure the holistic structures and work inferior for the images of weak texture. Furthermore, transformer-based methods [46, 58] with long-range dependency are utilized to firstly fill low-resolution tokens, and then upsample them with CNNs. Unfortunately, transformers demand huge memory footprint for large images. And the resolution disparity between transformer and CNN causes serious error propagation. On the other hand, some methods utilize auxiliary information for structure recovery, *e.g.*, edges [18, 38], segmentation [32, 43], and gradients [53]. Cao *et al.* [4] propose a sketch tensor space consisting of edges and wireframes [22] to facilitate holistic structure learning in man-made scenes. However, these sophisticated methods are usually based on multi-stage or multi-model designs, which are costly to be trained from scratch. Moreover, many researches [23, 33, 51] show that the position information is critical to learning the network, such as GANs [33, 51] and NeRF [36]. To our knowledge, there is no previous work that has explicitly discussed and utilized the position information in image inpainting.

Therefore, this motivates our work of incrementally inferring the holistic structural information and positional information to boost the performance of the inpainting model. Specifically, we leverage a transformer-based model to tackle holistic structures with edges and lines as the sketch tensor space [4]. Critically, such a normalized grayscale space can be easily upsampled by a simple CNN to higher resolutions without information loss. Further, we propose a novel incrementally training strategy with Zero-initialized Residual Addition (ZeroRA) [2] to incorporate the structural information into a pretrained inpainting model. This incremental strategy enjoys fast convergence for much fewer steps compared with retraining a new auxiliary-based model. Furthermore, we introduce the positional encoding for the mask region, which improves the performance of image restoration.

Formally, this paper proposes a novel ZeroRA based Incremental Transformer Structure (ZITS) inpainting framework enhanced with Masking Positional Encoding (MPE). Our ZITS has novel components of Masking Positional Encoding (MPE), Transformer Structure Restorer (TSR), Fourier CNN Texture Restoration (FTR), and Structure Feature Encoder (SFE). The TSR is composed of alternating axial [21] and standard attention blocks for the balance between the performance and the efficiency. Note that our TSR can achieve much better structure recovery compared with CNNs [4, 38]. The output grayscale edges and lines are upsampled with a simple 4-layer CNN. Then, a gated convolutions [57] based SFE encodes features and transfers

them to a FFC based inpainting model called FTR with ZeroRA. Furthermore, we use MPE to express both distances and directions from unmasked regions to masked ones.

We highlight several contributions as follows. (1) We propose using a transformer to learn a normalized grayscale sketch tensor space for inpainting tasks. Such an attention-based model can learn substantially better holistic structures with long-range dependency. (2) The auxiliary information can be incrementally incorporated into a pretrained inpainting model without retraining. (3) A novel masking positional encoding is proposed to improve the generalization of the inpainting model for different masks. (4) Extensive experiments on several datasets, which include Places2 [63], ShanghaiTech [22], NYUDepthV2 [37], and MatterPort3D [5] reveal that our proposed model outperforms other state-of-the-art competitors.

## 2. Related Work

**Inpainting by Auxiliaries.** Auxiliary information such as edges [38, 53], segmentation maps [32, 43], and gradients [53] are shown very useful to inpainting. Specifically, EdgeConnect [38] utilizes edges to help inpainting images with certain structures. Guo *et al.* [18] propose a two-stream network for image inpainting, which models the structure constrained texture synthesis and texture-guided structure reconstruction in a coupled manner. SGE-Net [32] just iteratively updates the semantic segmentation maps and the corrupted image. Cao *et al.* [4] further propose learning a sketch tensor space, composed of edges and lines for inpainting man-made scenes. In our work, we also take edges and lines as our auxiliary information. Differs from [4], the transformer is leveraged to rebuild edges and lines in ZITS. As some preliminary investigations [6] have shown its excellent capability in modeling structural relationships for natural image synthesis. Besides, almost all auxiliary-based methods need extra input channels for more information, which makes them must be retrained from scratch to utilize these additional inputs sufficiently. In our paper, we propose a flexible and effective way to add structural information to a pretrained inpainting model incrementally.

**Transformers for Image Generation.** Transformer [1, 45] achieved good performance on many tasks in the NLP and CV communities, as it learned long-range interactions on sequential data. Dosovitskiy *et al.* [11] firstly propose the use of a transformer for image recognition and show its great capacity. Many works [6, 29, 39] devote to reducing the time and space complexity for transformers. Esser *et al.* [13] and Ramesh *et al.* [41] leverage discrete representation learning for lower computational cost. The transformer is also used in image inpainting [46, 58] for low-resolution images reconstruction, and then guides the GAN-based CNN for further high-quality results. In our work, the transformer is used to build the holistic structure reconstruction

tion and then to guide the image inpainting, which enjoys excellent performance compared with CNN-based methods.

### 3. Method

**Overview.** The whole pipeline of ZITS is shown in Fig. 2. Given masked image  $\mathbf{I}_m$ , canny edge  $\mathbf{I}_e$  [10], lines  $\mathbf{I}_l$  [4], and binary mask  $\mathbf{M}$ , we concatenate and input them to the Transformer Structure Restoration (TSR) model for recovered edges and lines as sketch space  $[\tilde{\mathbf{I}}_e, \tilde{\mathbf{I}}_l] = \text{TSR}(\mathbf{I}_m, \mathbf{I}_e, \mathbf{I}_l, \mathbf{M})$  (Sec. 3.1). During the inference stage, the Simple Structure Upsampler (SSU) can easily upsample the grayscale sketch maps into arbitrary sizes (Sec. 3.2). Then, a gated convolution based Structure Feature Encoder (SFE) extracts multi-scale features  $\mathbf{S}_k = \text{SFE}(\tilde{\mathbf{I}}_e, \tilde{\mathbf{I}}_l, \mathbf{M}), \{k = 0, 1, 2, 3\}$  from upsampled sketches. We incrementally add  $\mathbf{S}_k$  to related layers of the Fourier convolution based CNN Texture Restoration (FTR) as  $\tilde{\mathbf{I}} = \text{FTR}(\mathbf{I}_m, \mathbf{M}, \alpha_k \cdot \mathbf{S}_k), \{k = 0, 1, 2, 3\}$  with the residual addition of zero-initialized trainable coefficients  $\alpha_k, i.e.,$  ZeroRA (Sec. 3.3).

#### 3.1. Transformer Structure Restoration

Since the transformer shows an ability to get expressive global structure recoveries [46], we leverage the capacity of the transformer to restore holistic structures in a relatively low resolution. For the input masked image  $\mathbf{I}_m$ , edges  $\mathbf{I}_e$ , lines  $\mathbf{I}_l$ , and mask  $\mathbf{M}$  in  $256 \times 256$ , we firstly downsample them with three convolutions to reduce computation for the attention learning. Such simple convolutions can also inject beneficial convolutional inductive bias for vision transformers compared with the patch based MLP embedding [49]. Then we add a learnable absolute position embedding to the feature at each spatial position and get  $\mathbf{X} \in \mathbb{R}^{h \times w \times c}$  for the input to attention layers, where  $h, w = 32$  are height and width, and  $c = 256$  is the feature channel.

To overcome the quadratic complexity of the standard self-attention [45], we alternately use axial attention modules [21] and standard attention modules as shown in the top left of Fig. 2. The axial attention module can be implemented easily by adjusting the tensor shape for row-wise and column-wise and then processing them with dot product-based self-attention respectively. To improve the spatial relation, we also provide relative position encoding (RPE) [40] for each axial-attention module. For the input feature  $\mathbf{X} \in \mathbb{R}^{h \times w \times c}$ , we suppose that  $\mathbf{x}_{ri,rj}, \mathbf{x}_{ci,cj} \in \mathbb{R}^c$  mean feature vectors of rows  $i, j$  and columns  $i, j$  of  $\mathbf{X}$ . Then the row and column-wise RPE based axial attention scores  $\mathbf{A}^{row}, \mathbf{A}^{col}$  can be written as

$$\begin{aligned} \mathbf{A}_{i,j}^{row} &= \mathbf{x}_{ri} \mathbf{W}_{rq} \mathbf{W}_{rk}^T \mathbf{x}_{rj}^T + R_{i,j}^{row}, \\ \mathbf{A}_{i,j}^{col} &= \mathbf{x}_{ci} \mathbf{W}_{cq} \mathbf{W}_{ck}^T \mathbf{x}_{cj}^T + R_{i,j}^{col}, \end{aligned} \quad (1)$$

where  $\mathbf{W}_{rq}, \mathbf{W}_{rk}, \mathbf{W}_{cq}, \mathbf{W}_{ck}$  are trainable parameters for query and key in row and column;  $R_{i,j}^{row}$  is the trainable

RPE value between row  $i$  and  $j$ , and  $R_{i,j}^{col}$  means the RPE value between columns  $i, j$ . Then, the attention scores are processed by the softmax operation. To stabilize the training, we use the pre-norm trick in [50]. Compared with the  $O(n^2)$  complexity of the standard self-attention, the axial attention only has  $O(2n^{\frac{3}{2}})$ , which allows us can handle more attention layers for a better capacity. Besides, we also remain some standard attention modules for learning the global correlation. Our ablation shows that this setting can improve the performance with the same memory cost.

After the encoding of stacked transformer blocks, features are upsampled by three transpose convolutions for outputting structures in  $256 \times 256$ . We use the binary cross-entropy (BCE) loss to optimize the predicted continuous sketch structures of edges  $\tilde{\mathbf{I}}_e$  and lines  $\tilde{\mathbf{I}}_l$  as

$$\mathcal{L}_e = \text{BCE}(\tilde{\mathbf{I}}_e, \hat{\mathbf{I}}_e), \quad \mathcal{L}_l = \text{BCE}(\tilde{\mathbf{I}}_l, \hat{\mathbf{I}}_l), \quad (2)$$

where  $\hat{\mathbf{I}}_e$  means the binary ground truth canny edges, and  $\hat{\mathbf{I}}_l$  indicates the antialiasing lines map got from the masking augmented wireframe detector from [4].

#### 3.2. Simple Structure Upsampler

To capture holistic structures for possible high-resolution images, we should upsample the generated edges and lines to arbitrary scales without obvious degeneration. However, vanilla interpolation-based resizing causes zigzag as shown in Fig. 3(f)–(i). Such artifacts are more serious for large image sizes, which deteriorate the inpainted results. Fortunately, the grayscale sketch tensor is easy to be upsampled with a learning-based method. At first we train a simple CNN as the SSU to upsample edges and lines to a doubled size. Although lines can be upsampled successfully, edges fail to get correct results as shown in Fig. 3(j). Because there are ambiguities in the canny edge from different image sizes as shown in Fig. 3(b) and Fig. 3(c). Since lines got from a wireframe parser have good discrete representations [22, 52], *i.e.,* a line can be indicated as positions of two end-points and their relation, we can draw line maps in various resolutions without any ambiguities as shown in Fig. 3(d) and Fig. 3(e). If the model is trained in lines, it can also achieve smooth high-resolution edge maps due to the generalization of the network as shown in Fig. 3(k). Through iteratively calling the SSU, we can get high-quality edges and lines with high resolutions.

#### 3.3. ZeroRA Structure Enhanced Inpainting

**Fourier CNN Texture Restoration (FTR).** For the texture restoration, we adopt the excellent work of [44] as our inpainting backbone. Suvorov *et al.* [44] propose to use Fourier convolutions [7] for the frequency domain learning, which can achieve resolution-robust inpainted results. As backbones used by other inpainting models [4, 38], FTR is an autoencoder model with several convolutions for down-sampling and up-sampling image features. The key module

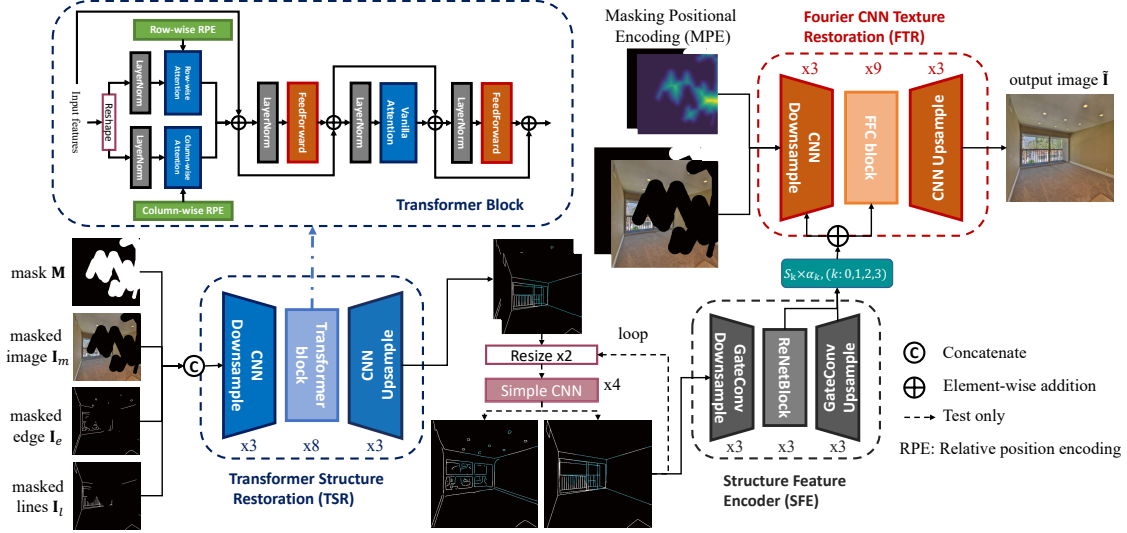


Figure 2. The overview of our ZITS. At first, the TSR model is used to restore structures with low resolutions. Then the simple CNN based upsampler is leveraged to upsample edge and line maps. Moreover, the upsampled sketch space is encoded by the SFE model, and added to the FTR through ZeroRA to restore the textures. The top left corner show details about the transformer block. The input feature are learned through row-wise and column-wise attentions respectively, then encoded by a standard attention module.

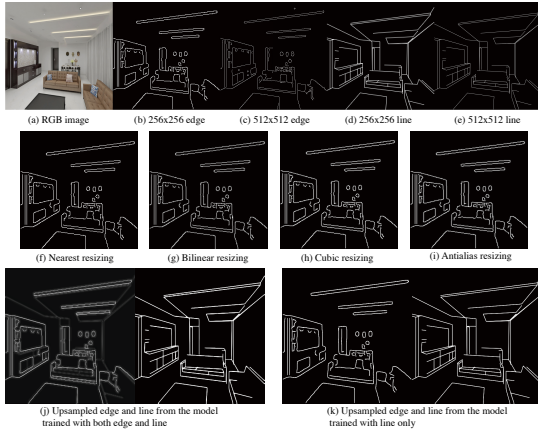


Figure 3. (a)–(e) indicate the ground truth images and structures. Edges are got from canny edge detector, while sigma is 2.0 for  $256 \times 256$  and 2.5 for  $512 \times 512$ . However, there are obvious ambiguities between (b) and (c). (f)–(i) show resizing edges of different interpolations. The learning based upsampling edge from (j) has significant superior quality compared with one from (k).

of FTR is the Fast Fourier Convolution (FFC) layer, which is consisted of two branches: 1) the local branch uses conventional convolutions and 2) the global branch convolutes features after the fast Fourier transform. Then two branches are combined for larger receptive fields and local invariance during the inpainting [44]. However, such a powerful model fails to learn reasonable holistic structures. And we further propose a series of novel components to improve it.

**Structure Feature Encoder (SFE).** For the given recovered edges  $\hat{I}_e$  and lines  $\hat{I}_l$  in arbitrary scales, we need a full convolutional network (FCN) to process them into a feature space. Our SFE is also an autoencoder model with 3 lay-

ers downsampling convolutions (encoder), 3 layers residual blocks with dilated convolutions [55] (middle), and 3 layers upsampling convolutions (decoder). For the encoder and the decoder in SFE, we use gated convolutions (GCs) [57] to transfer useful features selectively. GC learns another sigmoid activation with the same channels. Then the sigmoid features are multiplied to the convoluted ones as outputs. Although GCs are widely used in image inpainting for the better generalization to irregular masks, we use GCs to filter useful features to FTR. Because the grayscale sketch space is sparse, and not all features are necessary for the inpainting. Then, 4 coarse-to-fine feature maps  $S_k, k \in \{0, 1, 2, 3\}$  from the last middle layer and 3 decoder layers are selected to transfer structural features to FTR as

$$S_0, S_1, S_2, S_3 = \text{SFE}(\hat{I}_e, \hat{I}_l, M), \quad (3)$$

where  $M$  indicates the resized binary mask.

**Masking Positional Encoding (MPE).** Although the zero-padding in CNNs can provide some position information [23], it only contains information about spatial anchors [51]. Therefore, central generated regions from GANs tend to repeat meaningless artifacts without specific position encoding. When the image size is large, the effect of zero-padding will be further weakened, which causes more repeated artifacts [33] to generators.

During the inpainting, position information for unmasked regions is unnecessary, because the model always knows ground truth unmasked image regions. However, we think that position information is still critical for masked regions, especially when mask areas are large for high-resolution images. Limited by the receptive fields of CNNs, the model with large masks may lose the direction and po-

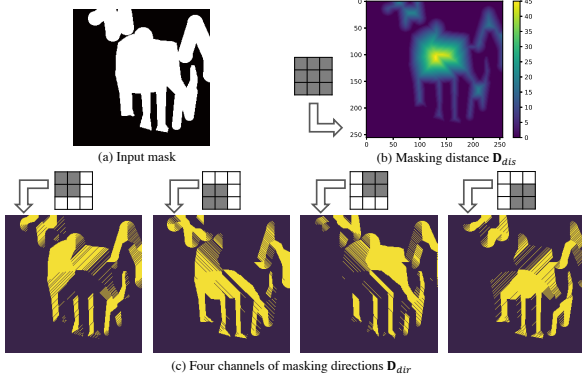


Figure 4. The illustration of our masking relative position encoding. (a) Input mask, (b) masking distance  $\mathbf{D}_{dis}$  and the all-one  $3 \times 3$  kernel, (c) masking directions  $\mathbf{D}_{dir}$  and their kernels.

sition information, which causes meaningless artifacts. Although FFC can extend the feature learning to the frequency domain, it is insensitive to distinguish masked or unmasked regions. Therefore, we propose to use position encoding in masked regions called MPE for the image inpainting, which is orthogonal to and improves upon the FFC in FTR.

Specifically, to clearly represent masked and unmasked positional relations with our MPE written as  $\mathbf{P}$ , it can be expressed as the masking distance  $\mathbf{P}_{dis}$  and the masking directions  $\mathbf{P}_{dir}$  as shown in Fig. 4. Given an inversed  $256 \times 256$  binary mask, where one indicates unmasked regions and zero indicates masked regions, we use a  $3 \times 3$  all-one kernel to calculate the masking distance  $\mathbf{D}_{dis}$  for each position in masked regions as shown in Fig. 4(b). Then, the distance is clipped and mapped by the Sinusoidal Positional Encoding (SPE) [45] to get  $\mathbf{P}_{dis} \in \mathbb{R}^{256 \times 256 \times d}$

$$\begin{aligned} \mathbf{P}_{dis,2i} &= \sin(\text{clip}(\mathbf{D}_{dis}, 0, D_{max})/10000^{\frac{i}{d}}), \\ \mathbf{P}_{dis,2i+1} &= \cos(\text{clip}(\mathbf{D}_{dis}, 0, D_{max})/10000^{\frac{i}{d}}), \end{aligned} \quad (4)$$

where  $i$  indicates the channel index;  $D_{max} = 128$ , and  $d = 64$  means the total channels of  $\mathbf{P}_{dis}$ , which is the same as the first convolution of FTR. Since SPE can only provide absolute positional information [51],  $\mathbf{P}_{dis}$  can be further resized by the nearest interpolation to various scales during the training for learning relative positional information in arbitrary resolutions. For masking directions, we use 4 different binary kernels to get the 4-channel one-hot vector  $\mathbf{D}_{dir} \in \mathbb{R}^{256 \times 256 \times 4}$ . Values of  $\mathbf{D}_{dir}$  depend on which kernel covers the masked regions firstly.  $\mathbf{D}_{dir}$  shows the nearest direction from a masked position to an unmasked one as shown in Fig. 4(c). Note that the masking direction is a multi-label vector, because a pixel may have more than one shortest direction. Then  $\mathbf{D}_{dir}$  is projected to a  $d$  dimension features with learnable embedding parameters  $\mathbf{W}_{dir} \in \mathbb{R}^{4 \times d}$  as

$$\mathbf{P}_{dir} = \mathbf{D}_{dir} \times \mathbf{W}_{dir} \in \mathbb{R}^{256 \times 256 \times d}. \quad (5)$$

$\mathbf{P}_{dis}$  and  $\mathbf{P}_{dir}$  are added as MPE to the first layer of FTR.

**Zero-initialized Residual Addition (ZeroRA).** Since most inpainting methods are based on sophisticated GANs nowadays, training the inpainting model incrementally is non-trivial. However, benefiting from various auxiliary information [4, 32, 38], incrementally training is flexible to improve the image inpainting. To improve the pretrained inpainting model incrementally with holistic structures, we propose to use ZeroRA, which has been leveraged in [2] to replace the layer normalization in the transformer. The idea of ZeroRA is simple. For the given input feature  $x$ , the output feature  $x'$  is got from adding a skip connection with function  $F$  to  $x$  with a zero-initialized trainable residual weight  $\alpha$  as

$$x' = x + \alpha \cdot F(x). \quad (6)$$

For simple linear-based models, if  $\alpha$  is initialized in zero, the input-output Jacobian will be initialized to 1, which makes the training stable. For more complex cases, experiments in [2] also prove the effectiveness of ZeroRA. Since ZeroRA can replace the layer normalization in the transformer, it can also improve the expressive power of the model without degrading variances to early layers.

In our case, we use ZeroRA to incrementally add structural information from SFE to FTR. Specifically, 4 zero-initialized  $\alpha_k, k \in \{0, 1, 2, 3\}$  are utilized to fuse 4 related feature maps  $\mathbf{S}_k$  from SFE. For the feature  $\mathbf{X}_k$  of FTR encoder layer  $k$ , which is based on Conv-BatchNorm-ReLU, we add residuals as follows

$$\begin{aligned} \mathbf{X}_{k+1} &= \text{Conv}(\mathbf{X}_k + \alpha_k \cdot \mathbf{S}_k), \\ \mathbf{X}_{k+1} &= \text{ReLU}(\text{BatchNorm}(\mathbf{X}_{k+1})). \end{aligned} \quad (7)$$

There is another advantage of the ZeroRA based incremental learning. The model output is equivalent to the pretrained one at the beginning of finetuning, which can effectively stable the training, and transfer necessary information adaptively. Our ablation studies show that the ZeroRA is important to incrementally finetune the pretrained inpainting model with additional information.

### 3.4. Loss Functions

We adopt the same loss functions as [44], which include L1 loss, adversarial loss, feature match loss, and high receptive field (HRF) perceptual loss [44]. Firstly, L1 loss is only calculated between the unmasked regions as

$$\mathcal{L}_{L1} = (1 - \mathbf{M}) \odot |\hat{\mathbf{I}} - \tilde{\mathbf{I}}|_1, \quad (8)$$

where  $\mathbf{M}$  indicates 0-1 mask that 1 means masked regions;  $\odot$  means the element-wise multiplication;  $\hat{\mathbf{I}}, \tilde{\mathbf{I}}$  indicate the ground truth and predicted images respectively. The adversarial loss is consisted of the discriminator loss  $\mathcal{L}_D$  and the generator loss  $\mathcal{L}_G$ . Moreover, we only regard features from masked regions as fake samples in  $\mathcal{L}_D$ . The PatchGAN [24] based discriminator is written as  $D$  and the combination of FTR and SFE can be seen as the generator  $G$ , Then the adversarial loss can be indicated as

$$\begin{aligned}
\mathcal{L}_D &= -\mathbb{E}_{\hat{\mathbf{I}}} \left[ \log D(\hat{\mathbf{I}}) \right] - \mathbb{E}_{\tilde{\mathbf{I}}, \mathbf{M}} \left[ \log D(\tilde{\mathbf{I}}) \odot (1 - \mathbf{M}) \right] \\
&\quad - \mathbb{E}_{\tilde{\mathbf{I}}, \mathbf{M}} \left[ \log(1 - D(\tilde{\mathbf{I}})) \odot \mathbf{M} \right], \\
\mathcal{L}_G &= -\mathbb{E}_{\hat{\mathbf{I}}} \left[ \log D(\hat{\mathbf{I}}) \right], \\
\mathcal{L}_{adv} &= \mathcal{L}_D + \mathcal{L}_G + \lambda_{GP} \mathcal{L}_{GP},
\end{aligned} \tag{9}$$

where  $\mathcal{L}_{GP} = \mathbb{E}_{\hat{\mathbf{I}}} \|\nabla_{\hat{\mathbf{I}}} D(\hat{\mathbf{I}})\|^2$  is the gradient penalty [16] and  $\lambda_{GP} = 1e - 3$ . We also use the feature match loss [47]  $\mathcal{L}_{fm}$ , which is based on L1 loss between discriminator features of true and fake samples.  $\mathcal{L}_{fm}$  is usually used to stable the GAN training. It can also slightly improve the performance. Furthermore, we use the HRF loss  $\mathcal{L}_{hrf}$  in [44] as

$$\mathcal{L}_{hrf} = \mathbb{E} \left[ \left( \phi_{hrf}(\hat{\mathbf{I}}) - \phi_{hrf}(\tilde{\mathbf{I}}) \right)^2 \right], \tag{10}$$

where  $\phi_{hrf}$  indicates a pretrained segmentation ResNet50 with dilated convolutions. As discussed in [44], using HRF loss instead of the perceptual loss [26] can improve the quality of the inpainting model. The final loss of our model in the incremental training can be written as

$$\mathcal{L}_{final} = \lambda_{L1} \mathcal{L}_{L1} + \lambda_{adv} \mathcal{L}_{adv} + \lambda_{fm} \mathcal{L}_{fm} + \lambda_{hrf} \mathcal{L}_{hrf}, \tag{11}$$

where  $\lambda_{L1} = 10$ ,  $\lambda_{adv} = 10$ ,  $\lambda_{fm} = 100$ ,  $\lambda_{hrf} = 30$ .

## 4. Experiments

### 4.1. Datasets

The proposed ZITS is trained on two datasets: Places2 [63] and our custom indoor dataset (Indoor). For Places2, we use about 1,800k images from various scenes as the training set, and 36,500 images as the validation. To better demonstrate the structural recovery, we collect 5,000 images from ShanghaiTech [22] and 15,055 images from NYUDepthV2 [37] to build the custom 20,055 Indoor training dataset. For the Indoor validation, we collect 1,000 images which are consist of 462 and 538 images from ShanghaiTech and NYUDepthV2 respectively. Places2 and Indoor can all be tested in both  $256 \times 256$  and  $512 \times 512$ . Besides, we also test the inpainting ability on high-resolution MatterPort3D [5] with 1,965 indoor images in  $1024 \times 1024$ . More details and results of MatterPort3D are discussed in the supplementary.

### 4.2. Implementation Details

**Training Settings.** Our ZITS is implemented with PyTorch. For the training of TSR, we use the Adam optimizer of learning rate  $6e-4$  with 1,000 steps warmup and cosine decay. TSR is trained with 150k and 400k steps for Indoor and Places2. On the other hand, we first train the FTR with Adam optimizer of learning rates  $1e-3$  and  $1e-4$  for generator and discriminator respectively. And FTR is trained with 100k steps on Indoor and 800k steps on Places2. Then, we incrementally finetune them with ZeroAR for just 50k steps

Table 1. Quantitative results on Indoor and Places2 in  $256 \times 256$ .

	Indoor				Places2			
	PSNR↑	SSIM↑	FID↓	LPIPS↓	PSNR↑	SSIM↑	FID↓	LPIPS↓
EC	24.07	0.884	22.02	0.135	23.31	0.839	6.21	0.149
MST	24.52	0.894	21.65	0.122	24.02	0.862	3.53	0.137
HiFill	-	-	-	-	20.76	0.770	21.33	0.246
Co-Mod	-	-	-	-	22.57	0.843	1.49	0.122
LaMa	25.20	0.902	16.97	0.112	24.37	0.869	1.63	0.155
Ours	<b>25.57</b>	<b>0.907</b>	<b>15.93</b>	<b>0.098</b>	<b>24.42</b>	<b>0.870</b>	<b>1.47</b>	<b>0.108</b>

on both Indoor and Places2, and reduce the generator learning rate to  $3e-4$ . Besides, we warmup the learning rate for training the SFE with 2,000 steps. For the training of TSR and FTR, input images are resized into  $256 \times 256$ . For the incremental finetuning, we separately train two versions of ZITS, which are the version trained in  $256 \times 256$  and the version trained in random size from 256 to 512. The second model can handle some situations with higher resolution inputs. And the MPE is also changed to a relative position encoding for the random size training.

**Mask Settings.** To tackle the real-world object removal task, we follow the mask setting from [4], which includes irregular masking brushes and COCO [34] segmentation masks with masking rates from 10% to 50%. Different from [4], we randomly combine irregular and segmentation masks with 20% to improve the learning difficulty.

### 4.3. Comparison Methods

We compare the proposed model with other state-of-the-art methods, which include Edge Connect (EC) [38], Contextual Residual Aggregation (HiFill) [54], Multi-scale Sketch Tensor inpainting (MST) [4], Co-Modulation GAN (Co-Mod) [62], and Large Mask inpainting (LaMa) [44]. All competitors are compared in the Places2. We also re-train EC, MST, and LaMa for the Indoor dataset to discuss the structure recovery. Note that the LaMa compared below are all trained with the same total steps as ZITS.

### 4.4. Quantitative Comparisons

**Inpainting Results.** In Tab. 1, we utilize PSNR, SSIM [48], FID [20], and LPIPS [61] to assess the performance of all compared methods on the Indoor and Places2 datasets in  $256 \times 256$  with mixed segmentation and irregular masks. More results with different masking rates are shown in the supplementary. For Indoor, our ZITS can achieve the best results on all metrics. While MST can get slightly better results compared with EC, which is benefited by the usage of lines. LaMa can get more acceptable FID and LPIPS while our ZITS can achieve significant improvements based on LaMa due to the seamlessly embedded structural information and positional encoding. Note that the gap between ZITS and MST is also caused by the quality gap of the structure recovery as discussed below. For Places2, HiFill fails to get good results with large masks, which may be caused by its limited model capacity. Note that Co-Mod has a low

Table 2. Quantitative Precision (P.), Recall (R.) and F1-score (F1) of edges and lines on Indoor and Places2.

		Edge			Line			Avg
		P.	R.	F1	P.	R.	F1	F1
Indoor	MST	23.79	26.87	21.36	43.67	51.95	37.77	33.73
	Ours	<b>37.34</b>	<b>34.25</b>	<b>35.10</b>	<b>53.60</b>	<b>66.23</b>	<b>58.35</b>	<b>46.72</b>
Places2	MST	22.54	18.29	20.19	34.22	49.21	37.09	28.64
	Ours	<b>35.64</b>	<b>27.92</b>	<b>30.39</b>	<b>43.70</b>	<b>60.54</b>	<b>49.35</b>	<b>39.87</b>

Table 3. Ablation studies of MPE on 512×512 Places2 finetuned with dynamic resolutions from 256 to 512.

	PSNR↑	SSIM↑	FID↓	LPIPS↓
with MPE	<b>24.23</b>	<b>0.881</b>	<b>26.08</b>	<b>0.133</b>
w/o. MPE	24.20	0.880	26.29	0.135

FID and LPIPS on Places2. However, Co-Mod is trained with a sophisticated StyleGAN [27] with much more training data compared with others. And our ZITS can even achieve slightly better results compared with Co-Mod with limited data scale and training steps. In general, our method has superior performance compared with LaMa, which is valuable with only 50k finetune steps. And LaMa in Tab. 1 are trained with the same total steps as ZITS.

**Results of Edges and Lines.** We show quantitative results of edges and lines on Indoor and Places2 in Tab. 2. Our TSR can achieve much better results on both Indoor and Places2 compared with MST. It demonstrates that the transformer-based TSR is amenable to learning holistic structures in a sparse tensor space, which can benefit the results of ZITS a lot as shown in Tab. 1. Note that TSR results in Tab. 2 are based on Mask-Predict [8, 14, 17], which can enrich the structural generation by iteratively sampling outputs but does not improve the quantitative metrics. More about Mask-Predict are discussed in the supplementary.

#### 4.5. Qualitative Comparisons

**Inpainting Results.** We show the qualitative inpainting results of Indoor in Fig. 5 and Places2 in Fig. 6. Compared with other methods, our ZITS can tackle more reasonable structures, especially our method can obtain clearer borderlines. Furthermore, ZITS achieves prominent improvements in the structure recovery compared with LaMa. Note that both LaMa and ZITS are trained with the same steps.

**Results of Edges and Lines.** We compared the structure recovery results in Indoor of Fig. 7, which are compared between our transformer-based TSR and CNN-based model from MST. Our TSR can achieve more reasonable and expressive results of both edges and lines. More qualitative structural results are shown in the supplementary.

#### 4.6. Ablation Studies

Quantitative ablation studies on Indoor are shown in Tab. 4. MPE and GCs can slightly improve the performance of FTR. Besides, if adding structural information from TSR

Table 4. Ablation studies with different settings on Indoor.

FTR	SFE	MPE	ReZero	GateConv	PSNR↑	SSIM↑	FID↓	LPIPS↓
✓					25.20	0.902	16.97	0.112
✓		✓			25.31	0.903	16.44	0.110
✓	✓	✓		✓	25.28	0.905	16.15	0.102
✓	✓		✓	✓	25.46	0.906	16.22	0.107
✓	✓	✓	✓		25.51	0.906	16.15	0.103
✓	✓	✓	✓	✓	<b>25.57</b>	<b>0.907</b>	<b>15.93</b>	<b>0.098</b>

Table 5. Quantitative results of 512×512 in Indoor and Places2, and 1024×1024 MatterPort3D.

		PSNR↑	SSIM↑	FID↓	LPIPS↓
Indoor(512)	LaMa	24.42	0.911	21.48	0.826
	Ours	<b>25.36</b>	<b>0.919</b>	<b>18.76</b>	<b>0.823</b>
Places2(512)	HiFill	20.10	0.764	65.47	0.291
	Co-Mod	22.00	0.843	30.04	0.166
	LaMa	24.15	0.877	27.86	0.149
	Ours	<b>24.23</b>	<b>0.881</b>	<b>26.08</b>	<b>0.133</b>
MatterPort3D(1k)	LaMa	26.40	0.944	14.04	0.133
	Ours	<b>26.55</b>	<b>0.946</b>	<b>12.34</b>	<b>0.116</b>

without ZeroRA, the improvement is limited. So ZeroRA is useful for incremental learning with a good convergence. Moreover, the full model achieves the best performance.

**ZeroRA.** We also show line charts of PSNR and FID during the finetuning in Fig. 8 with and without ZeroRA. The blue curve without ZeroRA is unstable at the beginning of the finetuning, while the red one with ZeroRA enjoys better convergence and stability. Because adding extra structural features without ZeroRA leads to dramatic output changing, which harms the vulnerable GAN training.

**MPE.** We further exploit the effects of MPE in high-resolution inpainting. FTR is trained without MPE at first. Then we use the ZeroRA technique to finetune the model with and without MPE of the same steps. Results in Tab. 3 show that the simple MPE based finetuning effectively improves the 512-inpainting in FID. From Fig. 9, ZITS with MPE generates images with natural and smooth colors.

#### 4.7. Results of High-Resolution Inpainting

We also compare the results of HiFill, Co-Mod, LaMa, and our ZITS in Places2(512) in Tab. 5. Besides, LaMa and ZITS are further compared in Indoor(512) and MatterPort3D(1k) in Tab. 5. LaMa and ZITS are firstly trained in 256×256 and then finetuned with dynamic resolutions from 256 to 512 with 50k steps. Models tested in Indoor(512) and MatterPort3D(1k) are both trained in Indoor training set. For the Places2(512), we randomly select 1,000 samples from 36,500 for the 512 testing. Our ZITS can achieve prominent improvements compared with LaMa, which illustrates that our MPE and incremental structure enhanced training is effective for high-resolution inpainting. Besides, ZITS can also get better 1k results in MatterPort3D. More high-resolution results can be seen in the supplementary.



Figure 5. Qualitative results of Indoor dataset compared among EC [38], MST [4], LaMa [44], and ours. Zoom-in for details.

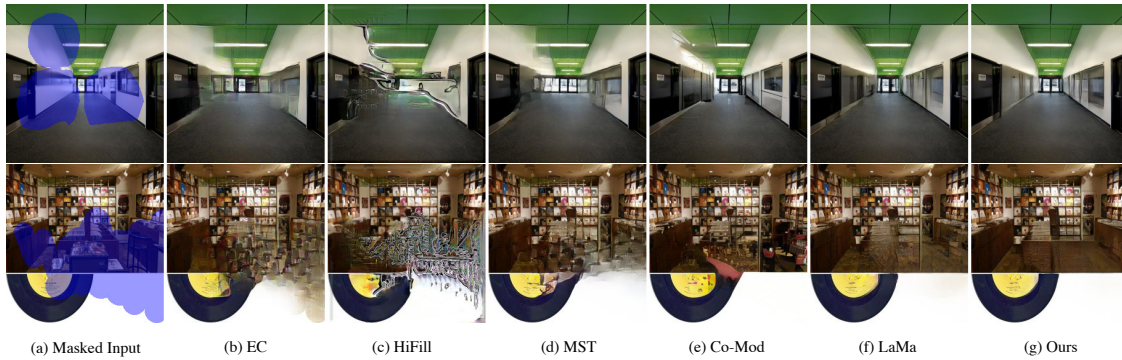


Figure 6. Qualitative results of Places2 compared among EC [38], HiFill [54], MST [4], Co-Mod [62], LaMa [44], and ours.

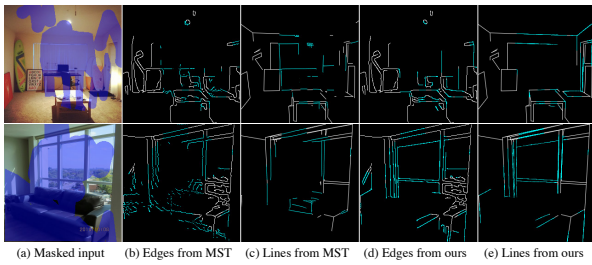


Figure 7. Qualitative results of edges and lines of Indoor dataset.

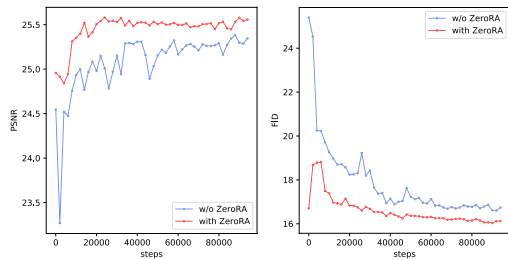


Figure 8. Structure enhanced finetuning with and without ZeroRA.

## 5. Conclusions

In this paper, we propose an incrementally structure enhanced inpainting model called ZITS. We use a transformer-based structure restorer to get much better holistic structures compared with previous methods. Then,



Figure 9. Ablations of  $512 \times 512$  Places2 with and without MPE.

a novel ZeroRA strategy is leveraged to incorporate auxiliary structures into a pretrained inpainting model with a few finetuning steps. The proposed masking positional encoding can further improve the inpainting performance. ZITS can achieve significant improvements based on the state-of-the-art model in experiments of various resolutions.

**Acknowledgement.** This work was supported by NSFC (62176061), SMSTC (2018SHZDZX01), and Shanghai Research and Innovation Functional Program (17DZ2260900).



## References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [2](#)
- [2] Thomas Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garrison W Cottrell, and Julian McAuley. Rezero is all you need: Fast convergence at large depth. *arXiv preprint arXiv:2003.04887*, 2020. [2](#), [5](#)
- [3] Marcelo Bertalmío, Guillermo Sapiro, Vicent Caselles, and Coloma Ballester. Image inpainting. *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000. [1](#)
- [4] Chenjie Cao and Yanwei Fu. Learning a sketch tensor space for image inpainting of man-made scenes. *arXiv preprint arXiv:2103.15087*, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)
- [5] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. [2](#), [6](#)
- [6] 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. [2](#)
- [7] Lu Chi, Borui Jiang, and Yadong Mu. Fast fourier convolution. *Advances in Neural Information Processing Systems*, 33, 2020. [3](#)
- [8] Jaemin Cho, Jiasen Lu, Dustin Schwenk, Hannaneh Hajishirzi, and Aniruddha Kembhavi. X-lxmert: Paint, caption and answer questions with multi-modal transformers, 2020. [7](#)
- [9] Antonio Criminisi, Patrick Perez, and Kentaro Toyama. Object removal by exemplar-based inpainting. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–II. IEEE, 2003. [1](#)
- [10] Lijun Ding and Ardeshir Goshtasby. On the canny edge detector. *Pattern Recognition*, 34(3):721–725, 2001. [3](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [2](#)
- [12] Omar Elharrouss, Noor Almaadeed, Somaya Al-Maadeed, and Younes Akbari. Image inpainting: A review. *Neural Processing Letters*, 51(2):2007–2028, 2020. [1](#)
- [13] 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. [2](#)
- [14] Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. Mask-predict: Parallel decoding of conditional masked language models. *arXiv preprint arXiv:1904.09324*, 2019. [7](#)
- [15] 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. [1](#)
- [16] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *arXiv preprint arXiv:1704.00028*, 2017. [6](#)
- [17] Junliang Guo, Zhirui Zhang, Linli Xu, Hao-Ran Wei, Boxing Chen, and Enhong Chen. Incorporating bert into parallel sequence decoding with adapters. *arXiv preprint arXiv:2010.06138*, 2020. [7](#)
- [18] Xiefan Guo, Hongyu Yang, and Di Huang. Image inpainting via conditional texture and structure dual generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14134–14143, 2021. [1](#), [2](#)
- [19] James Hays and Alexei A Efros. Scene completion using millions of photographs. *ACM Transactions on Graphics (SIGGRAPH 2007)*, 26(3), 2007. [1](#)
- [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, 2018. [6](#)
- [21] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019. [2](#), [3](#)
- [22] Kun Huang, Yifan Wang, Zihan Zhou, Tianjiao Ding, Shenghua Gao, and Yi Ma. Learning to parse wireframes in images of man-made environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 626–635, 2018. [2](#), [3](#), [6](#)
- [23] Md Amirul Islam, Sen Jia, and Neil DB Bruce. How much position information do convolutional neural networks encode? *arXiv preprint arXiv:2001.08248*, 2020. [2](#), [4](#)
- [24] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. [5](#)
- [25] Youngjoo Jo and Jongyoul Park. Sc-fegan: Face editing generative adversarial network with user’s sketch and color. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1745–1753, 2019. [1](#)
- [26] 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. [6](#)
- [27] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. [1](#), [7](#)
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. [1](#)
- [29] Manoj Kumar, Dirk Weissenborn, and Nal Kalchbrenner. Colorization transformer, 2021. [2](#)

- [30] Avisek Lahiri, Arnav Kumar Jain, Sanskar Agrawal, Pabitra Mitra, and Prabir Kumar Biswas. Prior guided gan based semantic inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13696–13705, 2020. 1
- [31] Anat Levin, Assaf Zomet, and Yair Weiss. Learning how to inpaint from global image statistics. *Proceedings Ninth IEEE International Conference on Computer Vision*, pages 305–312 vol.1, 2003. 1
- [32] Liang Liao, Jing Xiao, Zheng Wang, Chia-Wen Lin, and Shin’ichi Satoh. Guidance and evaluation: Semantic-aware image inpainting for mixed scenes. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII 16*, pages 683–700. Springer, 2020. 2, 5
- [33] Chieh Hubert Lin, Hsin-Ying Lee, Yen-Chi Cheng, Sergey Tulyakov, and Ming-Hsuan Yang. Infinitygan: Towards infinite-pixel image synthesis, 2021. 2, 4
- [34] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [35] 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
- [36] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pages 14134–14143, 2020. 2
- [37] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 2, 6
- [38] 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*, 2019. 2, 3, 5, 6, 8
- [39] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*, 2018. 2
- [40] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019. 3
- [41] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation, 2021. 2
- [42] Stefan Roth and Michael J. Black. Fields of experts: a framework for learning image priors. *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, 2:860–867 vol. 2, 2005. 1
- [43] Yuhang Song, Chao Yang, Yeji Shen, Peng Wang, Qin Huang, and C-C Jay Kuo. Spg-net: Segmentation prediction and guidance network for image inpainting. *arXiv preprint arXiv:1805.03356*, 2018. 2
- [44] 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. *arXiv preprint arXiv:2109.07161*, 2021. 1, 2, 3, 4, 5, 6, 8
- [45] 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, 3, 5
- [46] Ziyu Wan, Jingbo Zhang, Dongdong Chen, and Jing Liao. High-fidelity pluralistic image completion with transformers. *arXiv preprint arXiv:2103.14031*, 2021. 1, 2, 3
- [47] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. 6
- [48] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [49] Tete Xiao, Piotr Dollar, Mannat Singh, Eric Mintun, Trevor Darrell, and Ross Girshick. Early convolutions help transformers see better. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [50] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tiejian Liu. On layer normalization in the transformer architecture. In *International Conference on Machine Learning*, pages 10524–10533. PMLR, 2020. 3
- [51] Rui Xu, Xintao Wang, Kai Chen, Bolei Zhou, and Chen Change Loy. Positional encoding as spatial inductive bias in gans, 2020. 2, 4, 5
- [52] Nan Xue, Tianfu Wu, Song Bai, Fudong Wang, Gui-Song Xia, Liangpei Zhang, and Philip HS Torr. Holistically-attracted wireframe parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2788–2797, 2020. 3
- [53] Jie Yang, Zhiqian Qi, and Yong Shi. Learning to incorporate structure knowledge for image inpainting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12605–12612, 2020. 2
- [54] Zili Yi, Qiang Tang, Shekoofeh Azizi, Daesik Jang, and Zhan Xu. Contextual residual aggregation for ultra high-resolution image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7508–7517, 2020. 1, 2, 6, 8
- [55] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015. 1, 4
- [56] 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. [2](#)
- [57] 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. [2](#), [4](#)
- [58] Yingchen Yu, Fangneng Zhan, Rongliang Wu, Jianxiong Pan, Kaiwen Cui, Shijian Lu, Feiying Ma, Xuansong Xie, and Chunyan Miao. Diverse image inpainting with bidirectional and autoregressive transformers. *arXiv preprint arXiv:2104.12335*, 2021. [2](#)
- [59] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Baining Guo. Learning pyramid-context encoder network for high-quality image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1486–1494, 2019. [1](#)
- [60] Yu Zeng, Zhe Lin, Jimei Yang, Jianming Zhang, Eli Shechtman, and Huchuan Lu. High-resolution image inpainting with iterative confidence feedback and guided upsampling. In *European Conference on Computer Vision*, pages 1–17. Springer, 2020. [2](#)
- [61] 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](#)
- [62] 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. *arXiv preprint arXiv:2103.10428*, 2021. [6](#), [8](#)
- [63] 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](#), [6](#)