

Keys to Better Image Inpainting: Structure and Texture Go Hand in Hand

Jitesh Jain^{1,2,3*†} Yuqian Zhou^{4*†} Ning Yu⁵ Humphrey Shi^{1,3}

¹SHI Lab @ University of Oregon ²IIT Roorkee
³Picsart AI Research (PAIR) ⁴Adobe Inc. ⁵Salesforce Research

<https://praeclarumjj3.github.io/fcf-inpainting/>

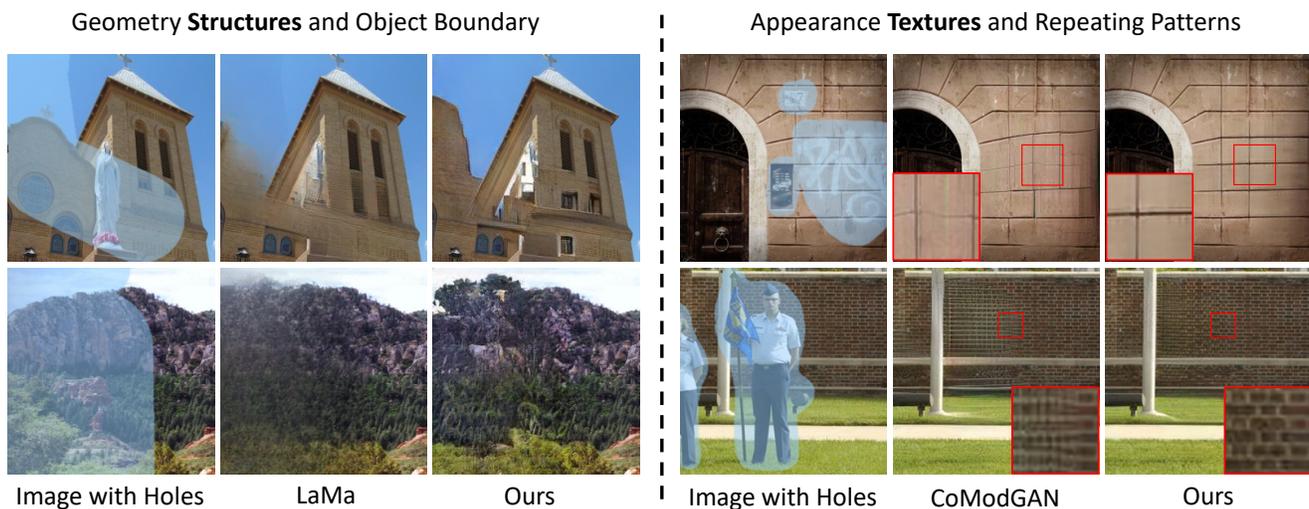


Figure 1: The most challenging issues for advanced image inpainting algorithms fall on generating better **structures** and **textures**. Left: LaMa [26] works well for repeating textures but generates fading out boundaries and structures when the holes get larger. Right: CoModGAN [45] with a StyleGAN-based [13] generator achieves impressive geometry structures, but it fails to reuse textures within the image to generate plausible repeating patterns. Our model generates good structures and textures simultaneously better than any state-of-the-arts.

Abstract

Deep image inpainting has made impressive progress with recent advances in image generation and processing algorithms. We claim that the performance of inpainting algorithms can be better judged by the generated structures and textures. Structures refer to the generated object boundary or novel geometric structures within the hole, while texture refers to high-frequency details, especially man-made repeating patterns filled inside the structural regions. We believe that better structures are usually obtained from a

coarse-to-fine GAN-based generator network while repeating patterns nowadays can be better modeled using state-of-the-art high-frequency fast fourier convolutional layers. In this paper, we propose a novel inpainting network combining the advantages of the two designs. Therefore, our model achieves a remarkable visual quality to match state-of-the-art performance in both structure generation and repeating texture synthesis using a single network. Extensive experiments demonstrate the effectiveness of the method, and our conclusions further highlight the two critical factors of image inpainting quality, structures, and textures, as the future design directions of inpainting networks.

* Equal Contribution.

† This work started when Jitesh interned at SHI Lab @ University of Oregon, and Yuqian was a Ph.D. student at IFP @ UIUC.

1. Introduction

Image inpainting aims to fill in missing parts of the incomplete input image such that an observer cannot distinguish between the inpainted regions and real regions of the output image. It has many applications in the industry, like object removal, photo retouching, and old photo restoration.

Traditionally, inpainting is achieved by diffusion-based [4] or patch-based methods [3]. They assume that the missing contents inside the hole regions can be synthesized by reusing textures or colors from the same image. These methods, especially patch-based ones, synthesize remarkable textures but mostly fail to complete semantic structures within the hole. GAN-based methods [37, 38] make the semantic structure generation possible. Among them, DeepFill [37, 38] first considered structure and texture synthesis. The model follows a two-stage network, where the first stage generates a coarse semantic map, and the second stage utilizes global contextual attention to copy similar deep features for texture enhancement. However, for most previous deep inpainting models [15, 34, 44, 52, 29, 17, 9, 40, 39, 53, 46], when the hole gets larger, the structure estimation becomes challenging. Recently, two milestones, LaMa [26], and CoModGAN [45] inspire us to study further the capability of deep networks handling inpainting structures and textures.

Zhao *et al.* [45] proposed the Co-Modulated Generative Adversarial Network (CoModGAN), which augments the encoded representation with a mapped stochastic noise vector to allow for stochasticity and better generation quality for images with large holes. The impressive image generation capability sources from the StyleGAN2-based [12, 13] generator, which follows a coarse-to-fine scheme. The conditional StyleGAN2 feeds the incomplete image’s global style and coarse structures and leverages rich structure data in the training dataset for novel generations. However, the generation quality relies more on the training data domain. Since CoModGAN does not include attention-related structures to enlarge the receptive field, the original image textures cannot be fully reused. CoModGAN performs poorly with novel textures and man-made repeating patterns.

The trend of image inpainting has been expected to be changed since LaMa [26] came out. Suvorov *et al.* [26] used the *Fast Fourier convolution* [6] inside their ResNet based LaMa-Fourier model to account for the lack of receptive field for generating repeating patterns in the hole regions. Before that, researchers struggled with global self-attention [38] and its high computational cost, but still cannot achieve reasonable recovery for repeating man-made structures as good as LaMa. Nevertheless, LaMa generates fading-out structures when the hole becomes larger and across object boundary. Recently, transformer-based methods [47, 28] model the global attention, while the structures can only be computed within a low-resolution

coarse image. Beyond that, good repeating textures cannot be synthesized. Recent diffusion-based inpainting models [24, 25, 19] pushed the limits of generative models, but the inference time can be too long for practical usage.

In this paper, we revisited the core design ideas of state-of-the-art deep inpainting networks. To address the issues mentioned above, we propose an intuitive and effective inpainting architecture that augments the powerful co-modulated StyleGAN2 [13] generator with the high receptiveness ability of FFC [6] to achieve equally good performance on both textures and structures as shown in Fig. 1. Specifically, we generate image structures in a coarse-to-fine StyleGAN-based generation scheme. Meanwhile, we merge between the generated coarse features and the skip features from the encoder and pass them through a Fast Fourier Synthesis (FaF-Syn) module to better generate repeating textures. Our idea is simple yet effective, synthesizing good structures and textures with a single network.

To summarize, we find that better structures can be obtained in a deeper coarse-to-fine GAN-based generator, and repeating textures are better synthesized with multi-scale high receptive field fourier convolutional layers. We combine the advantages of the two and propose a Fourier Coarse-to-Fine (FcF) generator for general-purpose image inpainting. Our model well handles textures and structures simultaneously and generalizes well to both natural and man-made scenes. Extensive experiments demonstrate the effectiveness of our proposed framework, and it achieves a new state-of-the-art on the CelebA-HQ dataset and remarkable performance comparable to state-of-the-arts on the Places2 dataset with a higher user preference rate.

2. Related Work

Traditional Image Inpainting. Traditionally, image inpainting tasks were resolved by either diffusion-based or exemplar-based methods. Diffusion-based methods use PDEs [4, 27, 31] or variational methods [2, 32] etc. to fill the hole by propagating the image pixels from the non-hole regions into the missing regions. It usually works well for connecting lines, curves within thinner holes since the smoothness constraints are used to regularize the hole-filling process, while for larger holes with more ambiguous structures, those approaches easily generate blur results. Copy-pasting of similar patches within the same images is categorized as the exemplar-based synthesis method. Both pixel-wise copying [7, 30] and patch-based synthesis [3, 33, 16] suffer from expensive nearest neighbor searching. Those methods produce good textures but messy structures.

Deep Image Inpainting. GAN-based [8] deep generative models [15, 34, 44, 52, 29, 17, 9, 40] were widely applied to inpainting tasks recently. Pathak *et al.* [23] first attempted to use GAN to address filling holes using seman-

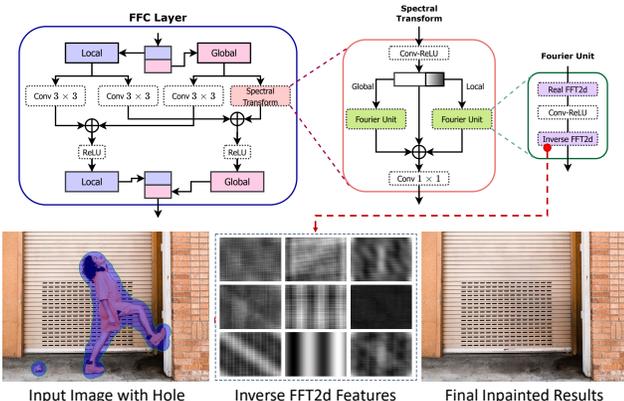


Figure 2: **FFC Layer and InverFFT2d feature visualization.** The FFC uses the Spectral Transform module in the global branch to account for the global context similar to LaMa [26]. The learned inverse FFT2d layer features explain why LaMa works well on repeating patterns. It actually generates global repeating patterns but not reconstructs image contents. The learned global repeating patterns are further merged inside the hole region to synthesize more complicated repeating patterns.

tically consistent contents. EdgeConnect [21] proposed to use edge detection results as the guidance for inpainting to form better structures. Later on, partial convolution [18] and gated convolution [37, 38] are proposed to tailor deep generative model for incomplete image feature extraction and reusing, making deep inpainting work for free-form holes. ProFill [42] then extended deepfillv2 [38] to apply iterative filling and confidence estimation to refine the textures. These methods fail to perform well on large irregular masks and textural images due to the small receptive field of the generator, lack of stochasticity, or larger memory and speed concerns like contextual attention [37]. Following the success of GAN-based methods, we formulate our inpainting framework based on the StyleGAN2 [13] architecture. The image generation ability associated with StyleGAN2 from stochasticity enables the large-hole filling with realistic structures.

Large Mask Inpainting. Recently, CoModGAN [45] proposed a co-modulation strategy using stochastic noise inside conditional StyleGAN2 [13] for improving image generation ability for large hole inpainting. Still, CoModGAN does not perform well when tested on texture-based images. For tackling repeating patterns in images, LaMa [26] proposed the use of Fast-Fourier Convolutions [6] inside the generator structure. However, LaMa produces a smooth and faded effect for large continuous masks. More recently, CMGAN [48] used FFC inside encoder and a cascaded global-spatial modulation-based decoder along with training on object-aware masks. However, CMGAN [48] struggles at good structure generation. In this work, we pro-

pose to combine the benefits of FFC and stochasticity using noise inside the co-modulated StyleGAN2 [13] coarse-to-fine generator to achieve robust performance on both textural and structural images for large free-form masks. The unification of stochasticity in co-modulated StyleGAN2 and FFC is non-trivial. It requires careful design of an architecture that does not collapse and behave effectively. Extensive experiments demonstrate that our integration prevents FFC from magnifying the noise in the coarse-level layers.

3. Methodology

In this section, we introduce the newly proposed network architecture, as shown in Fig. 3. The four-channel inputs concatenate the RGB masked image (I_{hole}) and the hole (M), where $I_{\text{hole}} = I_{\text{org}} \odot (1 - M)$. The inputs are fed into the encoder network (\mathcal{E}) to obtain the encoded latent vector z_{enc} and multi-level feature maps X_{skip} . Our generator network (\mathcal{G}) shares the spirit of the StyleGAN2 [13] architecture. Similar to CoModGAN [45], we generate random noise latent vector z and pass it through a mapping network (\mathcal{M}) to obtain the embedding z_w . z_w is concatenated with z_{enc} and fed into the generator \mathcal{G} . The core contribution is that we newly propose the Fast Fourier Synthesis Module (FaF-Syn) inside the Fourier Coarse-to-Fine (FcF) generator. More intuitions and details are introduced below.

3.1. Fourier Coarse-to-Fine (FcF) Generator

We aim to integrate the idea of LaMa, fast fourier convolutional residual blocks, into a co-modulated StyleGAN2-based coarse-to-fine generator. Intuitively, the coarse-to-fine generator renders global structures and image styles from the high-level feature and noise embedding. During the upsampling process in the generator, global texture features, both in the non-hole regions and in the generated hole regions, can be extracted by fast fourier convolutional layers and integrated appropriately to refine textures within the randomly generated structures. The idea is realized by a Fast Fourier Synthesis (FaF-Syn) module consisting of a Fast Fourier Residual (FaF-Res) Block. Within each FaF-Res block, there are two Fast Fourier Convolutional (FFC) layers. We will introduce them in a bottom-up order.

Fast Fourier Convolutional Residual Blocks (FaF-Res). The FaF Residual block in Fig. 3 (c) consists of two Fast Fourier Convolutional (FFC) layers (Fig. 2). The FFC [6] layer is based on a channel-wise fast Fourier transform (FFT) [5]. It splits channels into two branches: a) local branch uses conventional convolutions to capture the spatial details, and b) global branch uses a Spectral Transform module to consider the global structure and capture the long-range context. Finally, the outputs of the local and global branches are stacked together.

The Spectral Transform uses two Fourier Units (FU) to capture the global and semi-global information. The left

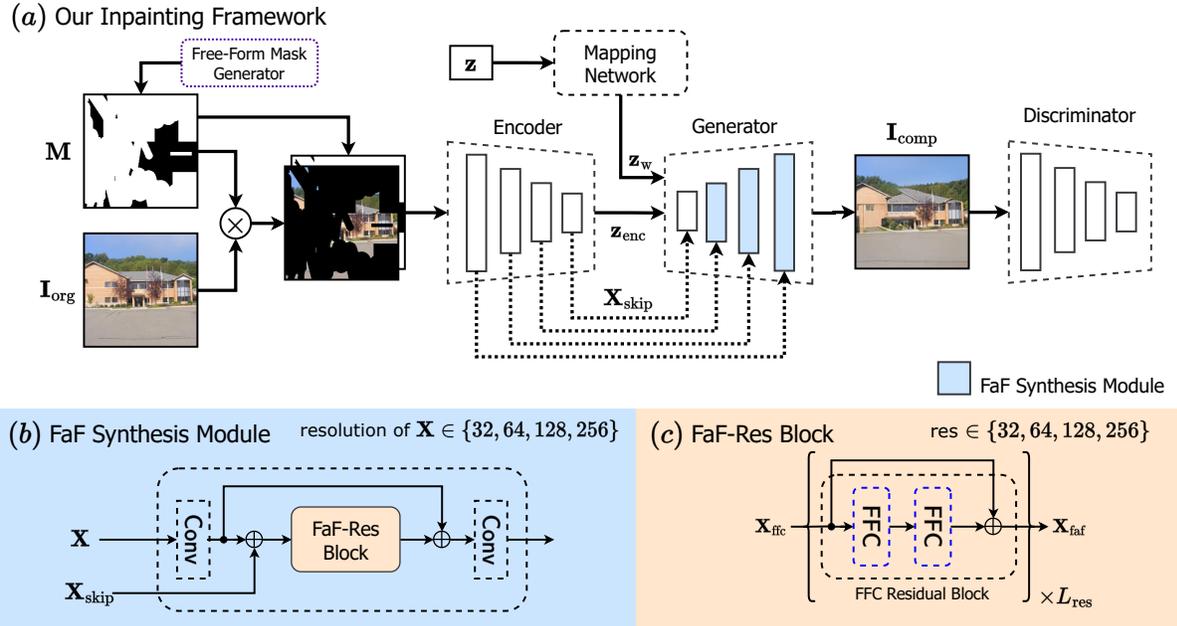


Figure 3: **Our model architecture.** (a) The inpainting framework. (b) The architecture of our FaF Synthesis (FaF-Syn) module inside the generator for resolutions $\in [32, 64, 128, 256]$. The convolutional layers inside FaF-Syn are co-modulated using the encoded features and style mapping of the latent noise vector. (c) The architecture of our FaF-Res Block.

Fourier Unit (FU) models the global context. On the other hand, the Local Fourier Unit (LFU) on the right side takes in a quarter of the channels and focuses on the semi-global information in the image. A Fourier Unit mainly breaks down the spatial structure into image frequencies using a Real FFT2D operation, a convolution operation in the frequency domain and finally recovering the structure using an Inverse FFT2D operation.

LaMa first applied FFC layers in inpainting yet did not reveal the reasons why it works for successfully synthesizing repeating patterns. We analyze LaMa’s intermediate features within the FFC layers and find that, after the inverse FFT2D layer within the Fourier Unit, the learned features do not represent and reconstruct complicated image contents directly, but generating multiple global repeating patterns, as shown in Fig. 2. The learned global repeating patterns are then merged inside the hole region to synthesize more complicated repeating contents. Therefore, in order to use FFC more effectively for inpainting, it is better to integrate the FFC layers into the generation process instead of feature encoding. It inspires us to carefully design a multi-scale FFC synthesis block and incorporate the FFC layers into the coarse-to-fine generator parts of the StyleGAN2.

Fast Fourier Synthesis (FaF-Syn) Module. Our generator (\mathcal{G}) shares a similar idea with CoModGAN [45], but the main difference is that we design the newly proposed Fast Fourier Synthesis (FaF-Syn) module (Fig. 3(b)) inside the coarse-to-fine generation process.

Integrating it into a StyleGAN2-based generator is not

trivial. There are two main problems to consider: First, the global repeating textures can be better modeled from the encoding features or in the generated features via a skip connection. Should we embed the FFC blocks into the encoder or the generator? We assume that it’s better to utilize it in the generation process by visualizing and analyzing the FFC features. Second, suppose we integrate the FFC blocks into the generator, the FFC layers may magnify the noisy generated structures in the very coarse level layers, causing unstable training and harming the performance. Which level of features will be better to include the FFC layers?

We empirically formulate our network in the following way: First, we use skip connections between \mathcal{E} and \mathcal{G} layers corresponding to the same resolution scale. Second, we introduce this Fast Fourier Synthesis (FaF-Syn) module as shown in Figure 3. FaF-Syn takes in both the encoded skip connected features X_{skip} , and the features X_{skip} upsampled from the previous level in the generator. FaF-Syn explicitly integrates the features from the encoder (*i.e.* existing image textures) and the generator (*i.e.* generated textures from the previous layers) to synthesize the global repeating textural features. It allows us to take advantage of the previous coarse-level repetitive textures and further refine them at the finer level. FaF-Syn is only applied to feature resolutions of 32×32 , 64×64 , 128×128 , and 256×256 . Our experiments show that applying it to a coarse level (like 8×8 and 16×16) harms the performance (Supplementary Material).

3.2. Other Modules

Encoder Network. Our encoder (\mathcal{E}) follows a similar architecture to the discriminator used in StyleGAN2 [13] but without the residual skip connections. \mathcal{E} takes \mathbf{I}_{hole} and \mathbf{M} downsamples it to a spatial size of 4×4 . We also use skip connections between \mathcal{E} and \mathcal{G} . Finally, we pass the flattened 4×4 encoded feature map through a linear layer to obtain an encoded latent vector \mathbf{z}_{enc} .

Mapping Network. We use a mapping network (\mathcal{M}) in our framework to transform our noise latent vector ($\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$) to an a latent space $\mathbf{z}_w = \mathcal{M}(\mathbf{z})$ [13]. We further perform an affine transform (\mathcal{A}) on a concatenation of \mathbf{z}_w and \mathbf{z}_{enc} from \mathcal{E} as $\mathbf{s} = \mathcal{A}(\text{stack}(\mathbf{z}_{\text{enc}}, \mathbf{z}_w))$. The style coefficient (\mathbf{s}) from \mathcal{A} is used to scale the weights of the convolutional layers inside our generator (\mathcal{G}). The architecture of \mathcal{M} is similar to the 8-layer MLP mapping network used in StyleGAN2 [13].

Discriminator. For our discriminator, we stick to the residual discriminator proposed in StyleGAN2 [13]. Our discriminator takes in a concatenation of hole masks \mathbf{M} and the original image \mathbf{I}_{org} or the completed image \mathbf{I}_{comp} depending on the training phase.

3.3. Loss Functions

We utilize a non-saturating logistic loss [8] with an R1-regularization [20] for our adversarial losses. We also use a reconstruction loss along with a high receptive field perceptual loss [26] to supervise the structures in the images during training. We find that reconstruction loss is important for learning the repeating patterns using FFC and the proposed FaF-Syn module.

Adversarial Loss. In the similar spirit of [13], we use the non-saturating cross-entropy loss for the adversarial training of our inpainting framework. The input of the discriminator is the concatenation of \mathbf{M} and real \mathbf{I}_{org} or fake \mathbf{I}_{comp} .

High Receptive Field Perceptual Loss. For the loss of the generator, similar to LaMa, we use a high receptive field perceptual loss (HRFPL) [26] which computes the ℓ_2 distance between \mathbf{I}_{comp} and \mathbf{I}_{org} , after mapping these images onto higher level features. The feature extractor is based on dilated ResNet-50 [35, 36] and is pre-trained for ADE20K [50, 51] semantic segmentation. Similar to [18], the loss can be represented as $\mathcal{L}_{\text{HRFPL}} = \sum_{p=0}^{P-1} \frac{\|\Psi_p^{\mathbf{I}_{\text{comp}}} - \Psi_p^{\mathbf{I}_{\text{org}}}\|_2}{N}$, where $\Psi_p^{\mathbf{I}^*}$ is the feature map of the p^{th} layer given an input \mathbf{I}^* , where N is the number of feature points in $\Psi_p^{\mathbf{I}_{\text{org}}}$.

Total Loss. We also include a pixel-wise reconstruction ℓ_1 loss between \mathbf{I}_{comp} and \mathbf{I}_{org} : $\mathcal{L}_{\text{rec}} = \|\mathbf{I}_{\text{comp}} - \mathbf{I}_{\text{org}}\|_1$ When calculating the final loss for the discriminator, we use a gradient penalty: $\mathcal{L}_{\text{reg}} = \mathbb{E}_{\mathbf{I}_{\text{org}}, \mathbf{M}} \left[\|\nabla \mathcal{D}_\theta(\text{stack}(\mathbf{M}, \mathbf{I}_{\text{org}}))\|^2 \right]$. The final loss is $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{adv}} + \lambda_{\text{rec}} \mathcal{L}_{\text{rec}} + \lambda_{\text{HRFPL}} \mathcal{L}_{\text{HRFPL}}$. The generator and discriminator are trained adversarially.

We empirically set $\lambda_{\text{rec}} = 10$, $\lambda_{\text{HRFPL}} = 5$, and $\lambda_{\text{reg}} = 5$ to balance the order of magnitude of each loss term.

4. Experiments

4.1. Datasets and Evaluation Metrics

We trained separate models on Places2 and CelebA-HQ datasets. Places2 [49] is a commonly-used dataset containing 8 million training images. We tested our models using the validation set consisting of 36,500 images. CelebA-HQ [11] is a high-quality image dataset of human faces containing 30,000 images. We divided the dataset into a training set with 26,000 images, a validation set with 2,000 images, and a test set with 2,000 images. We followed previous works to use LPIPS [43] and FID [10] as the evaluation metrics. We also conducted a user study to evaluate the perceptual quality of our results in a more faithful way.

4.2. Implementation Details

Network Details. The encoder \mathcal{E} downscales the input to a spatial size of 4×4 , increasing the channel dimension by $\times 2$ at each downscaled resolution to a maximum of 512 channels. We set the dimension for latent noise vector \mathbf{z} as 512. We flattened the output from the encoder to a dimension of 1024 to obtain \mathbf{z}_{enc} . We set the values of the Number (L_{res}) of FFC Residual Blocks at different resolutions as $\{L_{32} : 1, L_{64} : 1, L_{128} : 1, L_{256} : 1\}$.

Training Settings. We developed our codebase in PyTorch [22]. We conducted image completion at 256×256 resolution on the Places2 [49] and CelebA-HQ [11]. We trained our framework and CoModGAN[†] (for fair comparison) for 25M images on both Places2 and CelebA-HQ. When training on Places2, we randomly cropped 256×256 patches from the high-resolution images during training. We resized the CelebA-HQ images to 256×256 , following LaMa [26]. We randomly generated free-form masks during training following the generation strategy used in CoModGAN [45]. We used the Adam [14] optimizer with the learning rate set to 0.001. We used a batch size of 128.

Baselines. We compared our method to various baselines including the milestone works LaMa-Fourier [26] and CoModGAN [45], a transformer-based work called TFill [47], recent papers addressing structures and textures CTSDG [9] and CR-Fill [41], and some older works DeepFill-v2 [38] as well as Edge-Connect [21], and other well-performed work like AOT-GAN [40]. For most models except for CoModGAN, we used the publicly available codebase and pre-trained models. For fair comparison, since the public CoModGAN [45] checkpoint cannot be tested on 256×256 resolution, we trained our own PyTorch [22] implementation¹ of CoModGAN[†] with a reconstruction loss and used that for evaluation.

¹For our CoModGAN[†] re-implementation, we build on top of the the

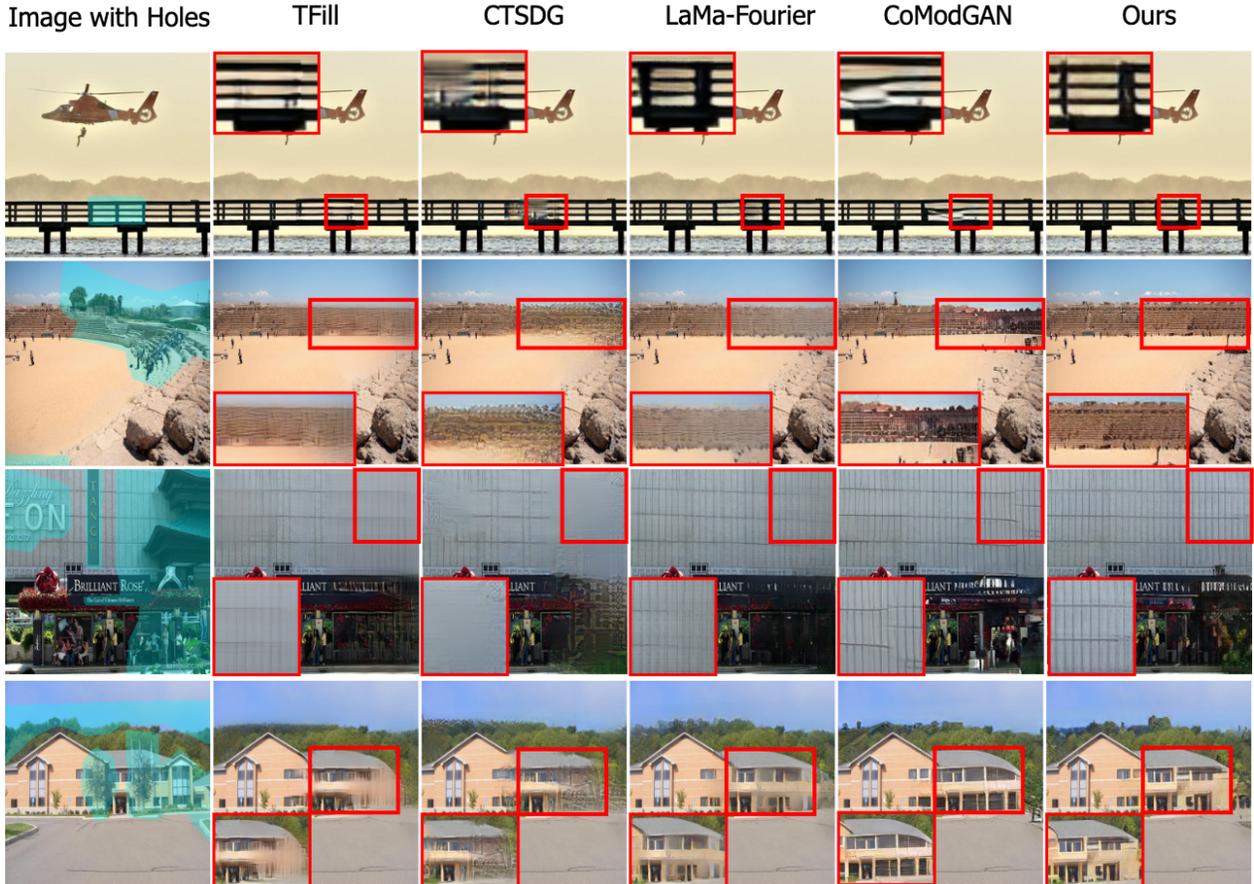


Figure 4: **Qualitative comparison to the state-of-the-art methods on Places2:** TFill [47], CTSDG [9], LaMa [26], CoModGAN[†] [45], and our framework (*Ours*). LaMa struggles to generate clear object boundaries while producing fading-out structures. CoModGAN does not have an attention scheme or large receptive field. Thus, it cannot effectively use self-similarity within the image and generates unseen and inconsistent textures. Ours handles structures and textures well in a single model. More results are in the supplementary material.

Evaluation Settings. When evaluating on Places2, we resized the images to 256×256 and tested with two different mask strategies: medium and segmentation [26] used in LaMa. Basically, the medium masks contain random strokes and rectangle boxes with medium size, and the segmentation masks were computed by replacing the masks of the segmentation onto other positions of the image. Please refer to LaMa [26] for more details. We used $30k$ and $4k$ samples for medium and segmentation masks respectively. We evaluated on CelebA-HQ for a total of $2k$ samples with medium and thick mask generation strategy [26].

4.3. Results and Comparisons

Qualitative Results. We compared the proposed FcF model to the highly relevant baselines including LaMa [26], CoModGAN[†] [45] (our PyTorch implementation), the lat-

StyleGAN2 [13] PyTorch codebase [link] as a more efficient alternative to the old version TensorFlow [1] codebase which is validated to be on par with the TensorFlow [1] code.

est transformer-based TFill [47] and the recent structure-texture inpainting network CTSDG [9]. The results on Places2 and CelebA-HQ are shown in Fig. 4 and 5.

As shown in Fig. 4, our model preserves much better repeating textures compared with CoModGAN. CoModGAN does not have any attention-related modules, so high-frequency features cannot be effectively reused given the limited receptive field. Our model enlarged the receptive field using fast Fourier layers and effectively rendered source textures on newly generated random structures. Meanwhile, ours also outperforms LaMa in generating object boundaries and structures. It is evident that LaMa generates fading-out artifacts when the hole reaches the image or object boundary. LaMa cannot hallucinate good structural information given large holes across longer pixel ranges. Ours, however, leverages the advantages of the coarse-to-fine generator to synthesize a clear shape boundary of objects in a better manner. In conclusion, our model integrates the advantages of two state-of-the-arts and simul-

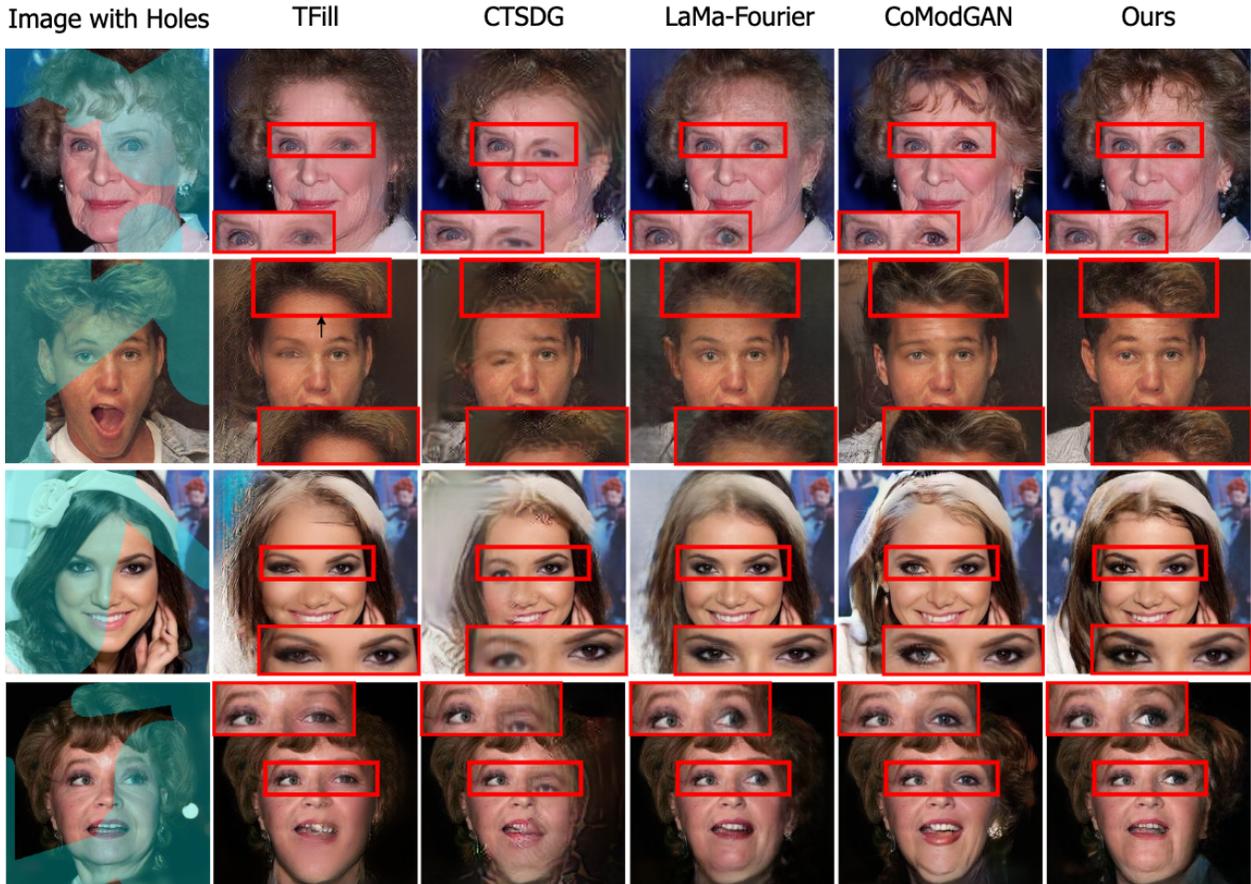


Figure 5: **Qualitative comparison to the state-of-the-art methods on CelebA-HQ dataset:** TFill [47], CTSDG [9], LaMa [26], CoModGAN[†] [45], and our framework *Ours*. The images are from the CelebA-HQ val (2k) dataset. LaMa mostly fades out the hair and generates a blurry boundary on the forehead. CoModGAN tends to generate unseen appearances inconsistent with the original face. Zoom in to check the eyes and eyebrows. Ours generate fine-detailed hairs and forehead shapes while preserving the original appearance of the person by generating consistent eyes and gaze direction. More results are in the supplementary material.

taneously generates remarkable structures and textures.

More qualitative evidence can be found in Fig. 5, which is more intuitive. While testing on face images, especially when we covered half of the faces, LaMa generates fading-out hairs on the forehead, and CoModGAN may use others’ eyes to complete the images. Though they both obtain good numbers in the quantitative results, some drawbacks are reflected, making both models not robust enough. Ours demonstrates a sound synthesis of hair and forehead shape and consistent eye and eyebrow appearance like LaMa. We can keep concluding that the proposed model work consistently well on both image structures and consistent textures.

Quantitative Results. We compared our method to several well-established baselines in Tab. 1. We found that LaMa and ours are always the top-two models and consistently outperform other baseline methods. Other baselines are not proven to work consistently well on larger masks.

CoModGAN is not working well on reconstruction. For Places2 evaluation, LaMa is still a strong baseline performing well in FID and the reconstruction-based metric LPIPS. Ours is comparable to the LaMa-Fourier model but is significantly better than the CoModGAN[†]. FFC layers and the proposed FaF-Syn modules add more global features to synthesize repeating textures for better background reconstruction. For the CelebA-HQ dataset, the proposed FcF model sets state-of-the-art while comparing with other baselines.

Due to the eco-friendly consideration, we include 256×256 resolution synthesis to prove concepts and draw scientific conclusions. In practice, we also trained a model on 512×512 resolution on Places2[49]. We use a batch size of 32 and the $\{L_{32} : 1, L_{64} : 1, L_{128} : 1, L_{256} : 1, L_{512} : 1\}$ setting while training. We achieve superior performance to the original CoModGAN [45] and are competitive to LaMa [26] as shown in Tab. 2. More qualitative comparison are in

Method	Places2 (256 × 256)				CelebA-HQ (256 × 256)			
	Medium Masks		Segm. Masks		Medium Masks		Thick Masks	
	FID↓	LPIPS↓	FID↓	LPIPS↓	FID↓	LPIPS↓	FID↓	LPIPS↓
Edge-Connect [21]	3.18	0.131	3.72	0.047	7.15	0.098	8.76	0.122
DeepFillv2 [38]	3.05	0.129	3.60	0.044	8.10	0.104	9.74	0.119
AOT-GAN [40]	1.95	0.116	3.31	0.043	8.27	0.104	13.89	0.135
CTSDG [9]	4.58	0.136	4.07	0.047	11.26	0.105	12.38	0.124
CR-Fill [41]	3.66	0.129	3.68	0.044	—	—	—	—
TFill [47]	2.52	0.120	3.24	0.042	6.49	0.090	6.54	0.102
CoModGAN [†] [45]	1.93	0.123	3.41	0.044	5.86	0.105	5.82	0.091
LaMa [26]	1.49	0.109	2.72	0.037	5.18	0.077	5.47	0.080
FcF (ours)	1.79	0.114	2.98	0.040	4.42	0.071	4.63	0.086

Table 1: **Quantitative evaluation on Places2 and CelebA-HQ.** We report LPIPS (↓) and FID (↓) metrics. The ↓ symbol means lower value signifies better performance. The **bold** text indicates the best performance, followed by **red** and **blue** fonts meaning the second and the third place.

Method	FID↓	LPIPS↓	User Preference (Baseline / Equal / Ours)
CoModGAN (official) [45]	2.32	0.045	21.33% / 17.33% / 61.33%
LaMa (official) [26]	2.00	0.040	39.33% / 12.00% / 48.67%
FcFGAN (ours)	2.06	0.041	- / - / -

Table 2: **Quantitative Comparisons using 512 × 512 images on Places2 [49]** for segmentation masks.

the Supplementary material. Thus, we demonstrate that our framework generalizes to higher resolutions equally well.

User Study. The existing metric LPIPS is hard to capture the enhanced textures and variant structures given complex scenes in Places2 [49]. FID is neither an ideal metric when we achieve equally good performance as LaMa [26] on man-made scenes in Places2 [49]. To further validate our model advantages, we conduct a user study via Amazon Mechanical Turk with 150 real user cases at 512 × 512 resolution. We let the users choose 'better', 'equal', or 'worse'. As shown in Tab. 2, our preference rate is the best, which further demonstrates our better visual quality.

4.4. Ablation Studies

Ablation on Number of FFC Residual Blocks. The number of FFC Residual Blocks inside our FaF-Res Block is an important tunable hyper-parameter. We experiment with various settings for $\{L_{32}, L_{64}, L_{128}, L_{256}\}$ in Sec. 4.4. We empirically observe that the setting $\{L_{32} : 1, L_{64} : 1, L_{128} : 1, L_{256} : 1\}$ gives the best performance.

Ablation on FaF-Syn Structures. We illustrated the inpainting results generated by different options of FaF-Syn module connections. We merged the encoder and decoder features in our current design before feeding them into the FaF-Syn Residual block. Alternatively, we experimented with two different ways: (1) directly connecting the FFC layers with the skipped features X_{skip} from the encoder (similar to using FFC inside encoder), or (2) connecting the skipped features with the FaF-Syn Residual block before merging to the generator feature X . The qualitative results

L_{32}	L_{64}	L_{128}	L_{256}	FID ↓	LPIPS ↓
0	0	0	0	13.53	0.275
0	1	1	1	12.14	0.266
0	1	2	2	11.92	0.263
0	2	2	2	12.77	0.268
1	1	1	1	11.33	0.264
2	2	2	2	15.33	0.280

Table 3: **Ablation on number of FFC Residual Blocks.** We find $\{L_{32} : 1, L_{64} : 1, L_{128} : 1, L_{256} : 1\}$ performs the best on the FID and LPIPS metrics.

Module	FID ↓	LPIPS ↓
FaF-Syn (ours)	11.33	0.264
FFC with X_{skip}	11.97	0.267
FaF-Res with X_{skip}	12.58	0.267
w.o. FFC	13.53	0.275

Table 4: **Ablation on Structures.** Our FaF-Syn performs the best on the FID and LPIPS metrics. The results show the effectiveness of the proposed design.

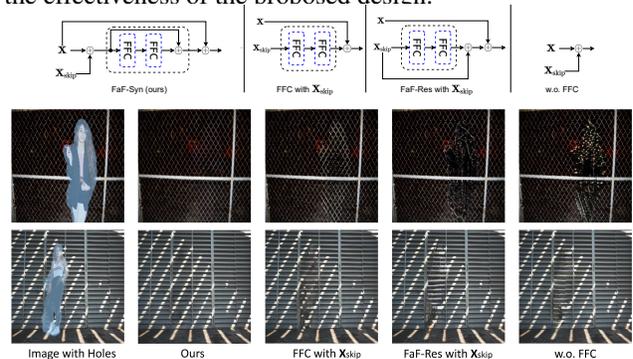


Figure 6: **Ablation study on alternatives of FaF-Syn module.** The results show the necessity of merging X before feeding X_{skip} into the FaF-Syn Residual block.

in Fig. 6 and quantitative comparison in Tab. 4 show the necessity of merging X and X_{skip} before feeding it into the FaF-Syn Residual block.

5. Conclusion

This work tackles the persistent challenges of synthesizing fair structures and textures in the hole regions. To this end, we propose a Fourier Coarse-to-Fine (FcF) inpainting framework that unites the receptive power of fast fourier convolutions to capture global repeating textures with the co-modulated coarse-to-fine generator to generate realistic image structures. Specifically, we proposed a simple yet effective FaF-Syn module aggregating the features from both the encoder and the generator to render textures on the generated structures progressively. Our model achieved a new state-of-the-art performance on the CelebA-HQ dataset and the best perceptual quality on the Places2 dataset. Extensive qualitative and quantitative analysis indicated that our framework is relatively robust to large masks and does not generate fading-out artifacts.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] Coloma Ballester, Marcelo Bertalmio, Vicent Caselles, Guillermo Sapiro, and Joan Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE TIP*, 2001.
- [3] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B. Goldman. Patchmatch: a randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 2009.
- [4] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *SIGGRAPH*, 2000.
- [5] E. O. Brigham and R. E. Morrow. The fast fourier transform. *IEEE Spectrum*, 1967.
- [6] Lu Chi, Borui Jiang, and Yadong Mu. Fast fourier convolution. In *NIPS*, 2020.
- [7] Alexei A Efros and Thomas K Leung. Texture synthesis by non-parametric sampling. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pages 1033–1038. IEEE, 1999.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [9] Xiefan Guo, Hongyu Yang, and Di Huang. Image inpainting via conditional texture and structure dual generation. In *ICCV*, 2021.
- [10] 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. *arXiv*, 2017.
- [11] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *ICLR*, 2018.
- [12] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [13] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020.
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [15] Avisek Lahiri, Arnav Kumar Jain, Sanskar Agrawal, Pabitra Mitra, and Prabir Kumar Biswas. Prior guided gan based semantic inpainting. In *CVPR*, 2020.
- [16] Lin Liang, Ce Liu, Ying-Qing Xu, Baining Guo, and Heung-Yeung Shum. Real-time texture synthesis by patch-based sampling. *ACM Transactions on Graphics (ToG)*, 20(3):127–150, 2001.
- [17] Liang Liao, Jing Xiao, Zheng Wang, Chia-wen Lin, and Shin’ichi Satoh. Guidance and evaluation: Semantic-aware image inpainting for mixed scenes. In *ECCV*, 2020.
- [18] 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 *ECCV*, 2018.
- [19] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022.
- [20] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. Which training methods for gans do actually converge? In *International Conference on Machine Learning (ICML)*, 2018.
- [21] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. Edgeconnect: Generative image inpainting with adversarial edge learning. In *ICCV*, 2019.
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [23] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.
- [24] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [25] Chitwan Saharia, William Chan, Huiwen Chang, Chris A Lee, Jonathan Ho, Tim Salimans, David J Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. *arXiv preprint arXiv:2111.05826*, 2021.
- [26] 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 *WACV*, 2022.
- [27] David Tschumperlé. Fast anisotropic smoothing of multi-valued images using curvature-preserving pde’s. *International Journal of Computer Vision*, 68(1):65–82, 2006.
- [28] Ziyu Wan, Jingbo Zhang, Dongdong Chen, and Jing Liao. High-fidelity pluralistic image completion with transformers. *arXiv preprint arXiv:2103.14031*, 2021.
- [29] Yi Wang, Ying-Cong Chen, Xin Tao, and Jiaya Jia. Vcnet: A robust approach to blind image inpainting. In *ECCV*, 2020.
- [30] Li-Yi Wei and Marc Levoy. Fast texture synthesis using tree-structured vector quantization. In *Proceedings of the*

27th annual conference on Computer graphics and interactive techniques, pages 479–488, 2000.

- [31] Joachim Weickert. Theoretical foundations of anisotropic diffusion in image processing. In *Theoretical foundations of computer vision*, pages 221–236. Springer, 1996.
- [32] Joachim Weickert. Theoretical foundations of anisotropic diffusion in image processing. In *Theoretical foundations of computer vision*, pages 221–236. Springer, 1996.
- [33] Zongben Xu and Jian Sun. Image inpainting by patch propagation using patch sparsity. *IEEE TIP*, 2010.
- [34] Zili Yi, Qiang Tang, Shekoofeh Azizi, Daesik Jang, and Zhan Xu. Contextual residual aggregation for ultra high-resolution image inpainting. In *CVPR*, 2020.
- [35] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016.
- [36] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *CVPR*, 2017.
- [37] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *CVPR*, 2018.
- [38] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *ICCV*, 2019.
- [39] ZHOU Yuqian, Elya Schechtman, Connelly Stuart Barnes, and Sohrab Amirghodsi. Image inpainting based on multiple image transformations, May 19 2022. US Patent App. 17/098,055.
- [40] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Bainig Guo. Aggregated contextual transformations for high-resolution image inpainting. *arXiv*, 2021.
- [41] Yu Zeng, Zhe Lin, Huchuan Lu, and Vishal M. Patel. Cr-fill: Generative image inpainting with auxiliary contextual reconstruction. In *ICCV*, 2021.
- [42] 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 *ECCV*, 2020.
- [43] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [44] 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 *CVPR*, 2020.
- [45] 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 *ICLR*, 2021.
- [46] Yunhan Zhao, Connelly Barnes, Yuqian Zhou, Eli Shechtman, Sohrab Amirghodsi, and Charless Fowlkes. Geofill: Reference-based image inpainting of scenes with complex geometry. *arXiv preprint arXiv:2201.08131*, 2022.
- [47] Chuanxia Zheng, Tat-Jen Cham, Jianfei Cai, and Dinh Phung. Bridging global context interactions for high-fidelity image completion. In *CVPR*, 2022.
- [48] Haitian Zheng, Zhe Lin, Jingwan Lu, Scott Cohen, Eli Shechtman, Connelly Barnes, Jianming Zhang, Ning Xu, Sohrab Amirghodsi, and Jiebo Luo. Cm-gan: Image inpainting with cascaded modulation gan and object-aware training. In *ECCV*, 2022.
- [49] 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*, 2017.
- [50] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017.
- [51] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *IJCV*, 2018.
- [52] Tong Zhou, Changxing Ding, Shaowen Lin, Xinchao Wang, and Dacheng Tao. Learning oracle attention for high-fidelity face completion. In *CVPR*, 2020.
- [53] Yuqian Zhou, Connelly Barnes, Eli Shechtman, and Sohrab Amirghodsi. Transfill: Reference-guided image inpainting by merging multiple color and spatial transformations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2266–2276, 2021.