

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# **GLaMa: Joint Spatial and Frequency Loss for General Image Inpainting**

Zeyu Lu<sup>†</sup>, Junjun Jiang<sup>\*†</sup>, Junqin Huang<sup>‡</sup>, Gang Wu<sup>†</sup>, Xianming Liu<sup>†</sup> <sup>†</sup> Harbin Institute of Technology <sup>‡</sup> Beihang University

### Abstract

The purpose of image inpainting is to recover scratches and damaged areas using context information from remaining parts. In recent years, thanks to the resurgence of convolutional neural networks (CNNs), image inpainting task has made great breakthroughs. However, most of the work consider insufficient types of mask, and their performance will drop dramatically when encountering unseen masks. To combat these challenges, we propose a simple yet general method to solve this problem based on the LaMa image inpainting framework [35], dubbed GLaMa. Our proposed GLaMa can better capture different types of missing information by using more types of masks. By incorporating more degraded images in the training phase, we can expect to enhance the robustness of the model with respect to various masks. In order to yield more reasonable results, we further introduce a frequency-based loss in addition to the traditional spatial reconstruction loss and adversarial loss. In particular, we introduce an effective reconstruction loss both in the spatial and frequency domain to reduce the chessboard effect and ripples in the reconstructed image. Extensive experiments demonstrate that our method can boost the performance over the original LaMa method for each type of mask on FFHQ [18], ImageNet [7], Places2 [42] and WikiArt [28] dataset. The proposed GLaMa was ranked first in terms of PSNR, LPIPS [39] and SSIM [34] in the NTIRE 2022 Image Inpainting Challenge Track 1 Unsupervised [27].

### 1. Introduction

Image inpainting, also known as image completion, has always been regarded as a challenge to fill the missing area of the image. Image inpainting can deal with various problems encountered in the real world, such as removing objects in photos, repairing damaged photos [32] or expanding photos. At the same time, image inpainting needs to maintain the coordination and semantic consistency between the repaired area and remaining parts of the image. Therefore, image inpainting also calls for strong generation ability.



Figure 1. Comparison with the LaMa [31], GLaMa and two-stage GLaMa for face Inpainting over some mask settings on FFHQ [18] dataset in  $1024 \times 1024$  resolution.

Nowadays, it has become a fundamental research topic in the field of computer vision and image processing society.

Thanks to the fast development of deep learning, the recovered results of these deep learning-based image inpainting approaches are getting better and better. Recent years, most state-of-the-art approaches are mainly based on convolutional neural networks or transformer. In the approaches of [22, 35, 38, 40], they apply the convolutional neural networks for image inpainting, while other line of research [33,37] leverages the transformer in image inpainting at the low-resolution image space, and then introduces the GAN based networks for high quality image generation. Suvorov et al. [31] utilize the Fast Fourier Convolution (FFC) instead of regular convolution to obtain features of global receptive fields in frequency domain. Most of the current work based on the semantic consistency (with the surrounding areas) can handle the "background completion" or "object removing" task very well.

Generally, these methods can effectively deal with some common image inpainting tasks. However, they still face some challenges. The trained model needs to be able to

<sup>\*</sup>Corresponding author (jiangjunjun@hit.edu.cn).



Figure 2. Overview of our proposed GLaMa.

deal with degraded image with various forms of mask, such as thin strokes, rectangle, and even extreme masks, because we do not know what image degradation process we will encounter. This is a challenging task because most of methods use specific masks for training, which may acquire poor results when meeting other masks that does not appear in the training processes. Recently, Andreas et al. [24] address the above issues. Some other diffusion models [29,30] can also deal with this challenge. However, it does not mean that the models based on convolutional neural networks (CNNs) or transformer cannot solve this problem well. In this work, we use more kinds of mask to enhance the robustness of the model. As shown in Fig. 4 and 5, our method can generate more realistic images.

When we visualize the spectra of real and fake images, we find that the difference between spectra of real pictures and fake images is obvious as shown in Fig. 3. It can be seen that there are many obvious errors in the spectra of the images generated by LaMa [31]. In addition, by looking closely, we can find that there will be many distinctive checkerboard effects and ripples on the images generated by LaMa [31]. Based on these findings, we add the reconstruction loss in the frequency domain as a regularization term to reduce the checkerboard effect and ripple of the generated image. We propose a joint spatial and frequency loss to train our network. As shown in Fig. 3, our method can do better in the frequency domain and produce more realistic results at the same time.

To demonstrate our superiority, we compare our method with state-of-the-art image inpainting approach [31] on multiple datasets. As shown in Tabs. 1 and 2, our method can boost the performance for each type of mask on FFHQ [18], ImageNet [7], Places2 [42] and WikiArt [28] datasets with the same network architecture and training epochs as LaMa [31]. And it was also ranked first on metric PSNR,

LPIPS [39] and SSIM [34] in the NTIRE 2022 Image Inpainting Challenge Track 1 Unsupervised [27]. The main contributions can be summarized as follows:

- We propose to explore the types of mask used in the training process. At the same time, using our mask generation strategy can effectively improve the results of the model.
- We use joint spatial and frequency loss in spatial domain and frequency domain with a regular term to reconstruct the image.
- As demonstrated in the experiments, our method can achieve significant performance improvements over LaMa [31] without changing the model architecture.

### 2. Related Work

#### 2.1. Image Inpainting

Early work of image inpainting [1-3, 6, 11] are modeldriven. They explore how to fill the missing information by exploiting clues from the local patch or neighbor patches in the input image.

Pathak et al. [26] propose the first deep learning image inpainting method that utilizes the CNN with an encoderdecoder architecture trained in the same way as GAN [10]. After that, a lot of methods based on CNN have been proposed. Iizuka et al. [14] improve the performance by exploiting a local-global discriminator, while Yu et al. [21] utilize a contextual attention model to model the long-distance context correlations. Due to the mask has negative influence on the results when using the regular convolution, several work modify the convolution operator, introducing partialconvolution [36] conv, gated-convolution [35]. [43] propose a mask awareness method using cascaded refinement network. Zeng et al. utilize "AOT Block" and "SoftGAN" to enhance the generator and discriminator [38]. A new GAN called "Co-Modulated GAN" combining conditional GAN and modulated GAN is introduced in [41].

There are also some work focusing on the fusion of local and global information. [22] utilizes feature equalization to fuse local and global features. Suvorov et al. [31] introduce the Fast Fourier Convolution (FFC) [5] and regular convolution to obtain features of global and local receptive fields in the frequency domain. These methods based on CNN could generate reasonable contents for masked regions.

In view of the dazzling performance of transformer architecture in other tasks, some transformer based methods have emerged recently. For example, Wan et al. [33] propose the first transformer based image inpainting method to get the image prior and send the image prior to a CNN. To incorporate the image prior, the approach of [37] designs a bidirectional and autoregressive transformer. More recently,



Figure 3. Visualization of the spectra of real and fake images. The spectra of images generated by our method is cleaner than the LaMa [31]. At the same time, the images of our method are more realistic and cleaner.

Dong et al. [8] design a transformer by edge auxiliaries to acquire prior and send the prior with masking positional encoding to a LaMa [31] like network. These methods based on transformer enjoy excellent performance compared with CNN-based methods.

In addition, there is another line of research based on diffusion model. SohlDickstein et al. [29] firstly utilize early diffusion models for image inpainting, while Song et al. [30] propose a score-based method using stochastic differential equations for image generation. More recently, Andreas et al. [24] propose a special model for image inpainting task which is based on diffusion model.

### 2.2. Loss Functions

In the image restoration task, the most commonly used reconstruction loss is  $L_1$  loss,  $L_2$  loss and Charbonnier loss. Some image restoration methods also use perceptual loss [16] and adversarial loss [10] to improve the perceptive performance. But above-mentioned methods all focus on the spatial domain reconstruction. In order to improve the reconstructed image quality, optimization in the frequency domain has gradually attracted researchers' attention in recent years. For example, spectral regularization is a preliminary attempt [23]. More recently, Gal et al. [9] propose a wavelet based image generation method. Jiang et al. [15] introduce the focal frequency loss which focuses on hard frequencies.Meanwhile, there are some other work [4, 17] on image restoration in the frequency domain.

## 3. Our Method

#### **3.1. Network Architecture**

Fast Fourier Convolution (FFC) in introduced in [5] to capture the global receptive field in the frequency domain. Inspired by this work, Suvorov et al. [31] design a network architecture called LaMa which achieves the state-of-the-art results using FFC to capture the global information. Following these pioneering work, we use LaMa as our network architecture. Like some other inpainting models, LaMa [31] uses an AutoEncoder model to extract the image features. The bottleneck module includes several FFC layers which are based on a channel-wise Fast Fourier transform (FFT), thus capturing global context information. The FFC layer splits feature channels into two branches: the local branch uses regular convolutions to obtain local features, while the global branch leverages FFC to obtain global features.

#### 3.2. Training with General Mask

Each training image x is from a training dataset superimposed by a synthetically generated mask. We train four different models for four datasets: FFHQ [18], ImageNet [7], Places2 [42] and WikiArt [28], respectively. After a lot of experiments, we find that the policy of mask generation noticeably influences the performance of the inpainting model as shown in Tab. 4.

We firstly try an aggressive large mask generation strategy based on LaMa [31]. This strategy uniformly uses samples from polygonal strip dilated by random strokes and rectangles of arbitrary aspect ratios. However, from the



Figure 4. Comparison with the state-of-the-art method: LaMa [31] for Face Inpainting over seven mask settings on FFHQ [18] dataset. The face images generated by our method are more realistic and have more details on the hair and facial features.

experimental results shown in Tab. 4, we can learn that the model trained with this mask generation strategy acquires poor results with the types of "Completion", "Expand", "Every n line" and "Nearest Neighbor" masks.

Subsequently, unlike the conventional practice, e.g. DeepFillv2 [35] or LaMa [31], we utilize seven types of mask generation strategies, "Completion", "Expand", "Every n line", "Nearest Neighbor", "Thin Strokes", "Medium Strokes", and "Thick Strokes", to generate samples randomly, which we call it *General Mask* generation strategy. These generated masks include not only small narrow masks, but also large-area masks. As shown in Tabs. 1 and 2, the model trained with *General Mask* generation strategy performs better than original LaMa [31] with all seven types of mask. This demonstrates that the diversity of degraded images can improve the robustness of the model to face various types of mask.

### 3.3. Loss Functions

The inpainting problem is ambiguous, which means one image that needs to be repaired might correspond to multiple images. One way is introduce some constraints to alleviate this. In this paper, we will introduce a joint spatial and frequency loss to regularize the optimization of our model, which regularizes the reconstruction results from both the spatial and frequency domains.

For the image inpainting task, the loss in the spatial domain is indispensable. First of all, we use  $L_1$  loss between the unmasked regions in the spatial domain. It can be formulated as:

$$\mathcal{L}_1 = \|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_1 \odot (1 - \boldsymbol{M}) \tag{1}$$

where x and  $\hat{x}$  indicate the ground truth and predicted images respectively. *M* represents 0-1 mask (1 means masked regions and 0 means no masked regions) and  $\odot$  means the element-wise multiplication. Moreover, we use the high receptive field perceptual loss  $\mathcal{L}_{PL}$ . It can be formulated as:

$$\mathcal{L}_{PL} = \left[\phi_{PL}(\boldsymbol{x}) - \phi_{PL}(\hat{\boldsymbol{x}})\right]^2 \tag{2}$$

where  $\phi$  indicates a pretrained segmentation ResNet50 [12] with dilated convolutions.

Denote  $\mathcal{L}_D$  the discriminator loss,  $\mathcal{L}_G$  the generator loss and  $\mathcal{L}_P$  the gradient penalty, the adversarial loss can be formulated as:

$$\mathcal{L}_{D} = -\mathbb{E}_{\boldsymbol{x}}[\log D(\boldsymbol{x})] - \mathbb{E}_{\hat{\boldsymbol{x}},\boldsymbol{M}}[\log D(\hat{\boldsymbol{x}}) \odot (1-\boldsymbol{M})] \\ -\mathbb{E}_{\hat{\boldsymbol{x}},\boldsymbol{M}}[\log(1-D(\hat{\boldsymbol{x}})) \odot \boldsymbol{M}]$$
(3)

Notably, we only regard features from masked regions as fake samples in  $\mathcal{L}_D$ .



Figure 5. Comparison with the state-of-the-art method: LaMa [31] for Inpainting over seven mask settings on ImageNet [7], Places2 [42] and WikiArt [28] datasets. The fourth column is sampled from ImageNet [7] dataset. The first, third and sixth columns are sampled from Places2 [42] dataset. The second, fifth and seventh columns are sampled from WikiArt [28] dataset.

$$\mathcal{L}_G = -\mathbb{E}_{\hat{\boldsymbol{x}}}[\log D(\hat{\boldsymbol{x}})] \tag{4}$$

$$\mathcal{L}_{P} = \mathbb{E}_{\boldsymbol{x}} \left\| \nabla_{\boldsymbol{x}} D(\boldsymbol{x}) \right\|^{2}$$
(5)

$$\mathcal{L}_{adv} = \mathcal{L}_D + \mathcal{L}_G + \lambda_P \mathcal{L}_P \tag{6}$$

We also use the feature match loss  $\mathcal{L}_{fm}$ , which is based on  $L_1$  loss between discriminator features of true and fake images, to stable the GAN training. Finally, the loss of LaMa [31] can be written as:

$$\mathcal{L}_{LaMa} = \lambda_1 \mathcal{L}_{L1} + \lambda_{adv} \mathcal{L}_{adv} + \lambda_{fm} \mathcal{L}_{fm} + \lambda_{PL} \mathcal{L}_{PL}$$
(7)

where  $\lambda_1 = 10$ ,  $\lambda_{adv} = 10$ ,  $\lambda_{PL} = 100$  and  $\lambda_{fm} = 30$ .

We find although LaMa [31] uses Fast Fourier Convolution, the model is not optimized in the frequency domain. Therefore, the spectrum of the images generated by LaMa [31] is defective, as shown in Fig. 3. With this in mind, we take the focal frequency loss [31] to construct the frequency fidelity term. Specifically, the focal frequency loss  $\mathcal{L}_{FFL}$  is defined as follows:

$$\mathcal{L}_{FFL} = \frac{1}{MN} \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} w(u,v) |F_r(u,v) - F_f(u,v)|^2$$
(8)

where  $F_r$  represents 2D discrete Fourier transform of real image (represents x here), and  $F_f$  represents 2D discrete

Fourier transform of fake image (represents  $\hat{x}$  here). The matrix w(u, v) represents the weight for the spatial frequency at coordinate (u, v), which is defined as:

$$w(u, v) = |F_r(u, v) - F_f(u, v)|^{\alpha}$$
(9)

where  $\alpha = 1$ . Furthermore, we find that the reconstructed images of LaMa [31] always get obvious checkerboard effect and ripples with the mask of "Nearest Neighbor" and "Every N Line". Based on this finding, we introduce the total variation loss, which is defined as: [29]:

$$\mathcal{L}_{TV} = \sum_{i,j} ((x_{i,j+1} - x_{i,j})^2 + (x_{i+1,j} - x_{i,j})^2)^{\frac{\beta}{2}} \quad (10)$$

where  $\beta = 2$ .

The final joint spatial and frequency loss function for our inpainting model can be written as:

$$\mathcal{L} = \alpha_1 \mathcal{L}_{TV} + \alpha_2 \mathcal{L}_{FFL} + \alpha_3 \mathcal{L}_{LaMa} \tag{11}$$

where  $\alpha_1 = 1$ ,  $\alpha_2 = 1$  and  $\alpha_3 = 1$ .

### 4. Experiments

### 4.1. Datasets and Metrics

In our experiments, we use FFHQ [18], ImageNet [7], Places2 [42] and WikiArt [28] datasets, which contain various kinds of images, to demonstrate the results. For FFHQ

Detecato	Mathada	LPIPS $(\downarrow)$						
Datasets	wienious	Completion	Expand	NearesNeighbor	ThinStrokes	EveryNLines	MediumStrokes	ThickStrokes
	LaMa	0.3302	0.6098	0.6511	0.0903	0.3125	0.1193	0.1374
Places2	Big LaMa	0.3298	0.6031	0.6473	0.0807	0.2788	0.1026	0.1266
	GLaMa	0.3212	0.5923	0.2280	0.0890	0.0891	0.1088	0.0774
	LaMa	0.3102	0.6089	0.6518	0.0976	0.312	0.1124	0.1307
FFHQ	Big LaMa	0.3048	0.5892	0.6081	0.1109	0.3067	0.1246	0.1331
	GLaMa	0.2932	0.5532	0.2402	0.1121	0.1829	0.1269	0.1324
	LaMa	0.3106	0.5877	0.6305	0.0768	0.2946	0.0997	0.1133
ImageNet	Big LaMa	0.3102	0.5834	0.6283	0.0607	0.2529	0.0815	0.1010
-	GLaMa	0.3009	0.5719	0.2061	0.0654	0.0651	0.0834	0.0859
	LaMa	0.3517	0.6261	0.6711	0.1093	0.3317	0.1397	0.1582
WikiArt	Big LaMa	0.3498	0.6217	0.6673	0.1008	0.2889	0.1219	0.1439
	GLaMa	0.3311	0.5989	0.2497	0.1065	0.0870	0.1244	0.0932

Table 1. LPIPS comparisons of different methods on the four datasets. The top two performing method is highlighted in red and blue.

Detecato	Mathada	FID $(\downarrow)$						
Datasets	wienious	Completion	Expand	NearestNeighbor	ThinStrokes	EveryNLines	MediumStrokes	ThickStrokes
	LaMa	33.63	87.35	163.92	6.37	17.07	10.17	13.21
Places2	Big LaMa	30.76	81.73	167.31	5.74	14.04	9.18	12.08
	GLaMa	29.98	68.74	8.16	6.29	2.05	10.02	13.03
	LaMa	23.82	117.60	133.70	6.60	21.44	8.64	9.19
FFHQ	Big LaMa	23.32	111.91	131.72	5.89	20.01	6.68	7.54
	GLaMa	20.37	89.70	7.65	6.10	6.21	7.65	8.09
	LaMa	25.28	119.34	135.01	8.69	23.90	10.53	11.19
ImageNet	Big LaMa	25.31	110.78	131.47	7.53	22.17	8.77	9.98
U U	GLaMa	22.91	91.78	9.78	8.06	8.21	9.24	10.01
	LaMa	39.13	92.01	162.59	12.88	23.47	16.75	18.96
WikiArt	Big LaMa	36.76	86.18	153.31	11.18	20.62	15.81	17.04
	GLaMa	35.72	74.51	12.08	11.89	8.03	16.23	17.87

Table 2. FID comparisons of different methods on the four datasets.

[18], we use about 60K face images as the training set and 2,000 other images as the validation. For Places2 [42], we use about 1,800K images from various scenes as the training set and 3650 other images from validation set as the validation. For ImageNet [7], we use about 1,000K images as the training set and 2000 other images from validation dataset as the validation set. And for WikiArt [28], we use the whole training set and 2000 other images from validation.

In the training stage for FFHQ [18], our images are resized to  $256 \times 256$  resolution. For the remaining three training datasets, we crop the images to  $256 \times 256$  pixels. Moreover, we use the original size of the images in the validation stage without cropping or resizing them.

We follow the established practice in recent image inpainting literature and use Learned Perceptual Image Patch Similarity (LPIPS) [39] and Frechet inception distance (FID) [13] metrics. Compared to PSNR and SSIM [34], LPIPS [39] and FID [13] are more suitable for measuring performance of inpainting for large masks.

#### 4.2. Implementation Details

For inpainting network [31], we followed LaMa [31] using a ResNet-like [12] architecture with 3 downsampling blocks, 6-18 residual blocks, and 3 upsampling blocks. We use Adam optimizer [20], with the fixed learning rates 0.001 and 0.0001 for inpainting and discriminator networks, respectively. We set the hyper-parameters using the grid search strategy on FFHQ dataset, which leds to the weight values  $\alpha_1 = 1$ ,  $\alpha_2 = 1$ ,  $\alpha_3 = 1$ ,  $\alpha = 1$ ,  $\beta = 2$ . The same parameters are used for other three datasets for training. All models are trained for 40 epochs with a batch size of 20, on 8 NVidia V100 GPUs for approximately 72 hours. The LaMa and Big LaMa [31] models used as baseline in the Tabs. 1 and 2 were retrained with the same experimental setups. Becaused LaMa and Big LaMa [31] and the results are higher than the original LaMa and Big LaMa [31] model in FFHQ [18] and Places2 [42] datasets.

#### 4.3. Comparisons to the Baselines

We compare the proposed approach with the strong baselines (LaMa and Big LaMa [31]). For each dataset, we validate the performance across seven types of mask. As shown in Tabs. 1 and 2, we can see that our GLaMa is better than the original LaMa [31] especially for "Completion" mask, "Expand" mask, "Nearest Neighbor" mask and "Every N Lines" mask. Compared with the images generated by LaMa, our images are more realistic with less checkerboard effect and ripples, as shown in Fig. 3, 4 and 5. GLaMa also works better than Big LaMa [31] in most types of mask with the same network architecture and training time as LaMa [31]. Meanwhile, our method that simply using the General Mask generation strategy and joint spatial and frequency loss can bring considerable improvements.

$\mathcal{L}_{LaMa}$	$\mathcal{L}_{TV}$	$\mathcal{L}_{FFL}$	FID $(\downarrow)$	LPIPS $(\downarrow)$
$\checkmark$			57.015	0.3382
$\checkmark$	$\checkmark$		22.363	0.2350
$\checkmark$	$\checkmark$	$\checkmark$	20.827	0.2253

Table 3. Ablation studies over the loss functions. Here, FID and LPIPS values of the results are reported.

LaMa Mask	LaMa Plus Mask	General Mask	$FID(\downarrow)$	LPIPS (4)
$\checkmark$			49.015	0.3118
	$\checkmark$		23.514	0.2448
		$\checkmark$	20.827	0.2253

Table 4. Ablation studies over the mask generation methods. *LaMa Mask, LaMa Plus Mask* represent using original LaMa, LaMa with "Nearest Neighbour" and "Every n line", respectively. *General Mask* represents using our method.

#### 4.4. Ablation Studies

The goal of the ablation studies is to carefully validate the influence of different components of our method. To investigate the effectiveness of ours mask generation strategy and losses, we conduct ablation experiments and the results comparison are shown in Tabs. 3 and 4. All these results are conducted on FFHQ [18] dataset.

Mask Generation Strategy. We verify that the General Mask generation strategy is necessary in Tab. 4. LaMa Mask and LaMa Plus Mask generation strategy represent using LaMa original policy to generate mask (three types of mask) and LaMa original policy with "Nearest Neighbour" and "Every n line" to generate mask (five types of mask), respectively. General Mask generation strategy represents using our policy to generate mask (seven types of mask). We can see from Tab. 4 that using General Mask generation strategy can greatly improve the model performance. This demonstrates that the diversity of degraded images can improve the robustness of the model to face various types of mask.

**Loss Functions.** We experiment with the two proposed losses in Tab. 3, where total variation loss [25] and focal frequency loss [15] can boost the model performance. At

the same time, it can be seen from Fig. 3 that the spectrum of the image generated by GLaMa is more accurate than LaMa [31] : ours spectrum is clean, while the spectrum of LaMa is very noisy. These experimental results further prove the superiority of our method.



Figure 6. Visual comparisons with unseen masks.

### 4.5. Generalization of the Proposed Method

**Generalization to the Masks.** To verify the generalization ability of our model for the masks that do not appear in the training stage, we randomly generate some other kinds of masks that are not in the seven kinds of masks. The mask generation strategy will combine the samples from seven kinds of masks and rectangular mask of arbitrary aspect ratios to generate some other kinds of masks. As shown in Fig. 6, GLaMa can generate high-quality images, while the images generated by LaMa [31] are low-quality, which proves that our method has strong generalization performance. Hence one can see that, whether the masks are appeared during the training stage, our model can recover the target images very well.

Generalization to the Resolution. To further exploit the reconstruction ability of different approaches, we also conduct some experiments when the input images are with  $512 \times 512$  pixels (or higher resolution) while the training images are with  $256 \times 256$  pixels. It can be seen from Figs. 1 and 7 that LaMa [31] is not good at repairing the large-area mask of high-resolution face. Some low-quality images are generated with the high-resolution images and large-scale masks. Since we use  $256 \times 256$  pixels can be repaired very well. Based on this finding, we design a two-stage model. In the first stage, we use the original LaMa [31] to generate the coarse image in  $256 \times 256$  resolution. In the



Figure 7. Some failed examples on the FFHQ [18] dataset.

Taama	FFHQ						
Teams	$FID(\downarrow)$	LPIPS $(\downarrow)$	PSNR (†)	SSIM (†)			
VIP	115.922	0.433	17.144	0.649			
AIIA (Ours)	9.823	0.239	25.316	0.814			
HSSLAB	13.504	0.236	25.187	0.821			
KwaiInpainting	21.345	0.239	25.060	0.838			
ArtificiallyInspired	4.719	0.205	25.999	0.816			
SIGMA	7.203	0.248	24.860	0.795			
Taama	Places2						
Teams	$FID(\downarrow)$	LPIPS $(\downarrow)$	PSNR (†)	SSIM (†)			
VIP	52.471	0.415	17.256	0.626			
AIIA (Ours)	8.772	0.224	24.145	0.800			
HSSLAB	9.861	0.227	24.345	0.798			
KwaiInpainting	18.334	0.255	23.410	0.787			
ArtificiallyInspired	7.544	0.225	23.248	0.777			
SIGMA	11.496	0.270	22.562	0.748			
	ImageNet						
Taoma		Imag	geNet				
Teams	FID (↓)	Imag LPIPS (↓)	geNet PSNR (†)	SSIM (†)			
Teams VIP	FID (↓) 50.898	Imag LPIPS (↓) 0.403	geNet PSNR (†) 17.450	SSIM (†) 0.626			
Teams VIP AIIA (Ours)	FID (↓) 50.898 10.007	Imaş LPIPS (↓) 0.403 0.207	geNet PSNR (↑) 17.450 25.226	SSIM (†) 0.626 0.800			
Teams VIP AIIA (Ours) HSSLAB	FID (↓) 50.898 10.007 11.770	Imag LPIPS (↓) 0.403 0.207 0.227	geNet PSNR (†) 17.450 25.226 24.303	SSIM (†) 0.626 0.800 0.798			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting	FID (↓) 50.898 10.007 11.770 18.854	Imag LPIPS (↓) 0.403 0.207 0.227 0.249	geNet PSNR (†) 17.450 25.226 24.303 23.804	SSIM (↑) 0.626 0.800 0.798 0.787			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired	FID (\$) 50.898 10.007 11.770 18.854 12.059	Imag LPIPS (↓) 0.403 0.207 0.227 0.249 0.217	geNet PSNR (†) 17.450 25.226 24.303 23.804 24.278	SSIM (†) 0.626 0.800 0.798 0.787 0.777			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired SIGMA	FID (\$) 50.898 10.007 11.770 18.854 12.059 19.646	Imag LPIPS (↓) 0.403 0.207 0.227 0.249 0.217 0.311	geNet PSNR (†) 17.450 25.226 24.303 23.804 24.278 22.454	SSIM (†) 0.626 0.800 0.798 0.787 0.777 0.748			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired SIGMA	FID (\$\) 50.898 10.007 11.770 18.854 12.059 19.646	Imag LPIPS (↓) 0.403 0.207 0.227 0.249 0.217 0.311 Wil	eeNet PSNR (†) 17.450 25.226 24.303 23.804 24.278 22.454 ciArt	SSIM (↑) 0.626 0.800 0.798 0.787 0.777 0.748			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired SIGMA Teams	FID (↓) 50.898 10.007 11.770 18.854 12.059 19.646 FID (↓)	Imag LPIPS (↓) 0.403 0.207 0.227 0.249 0.217 0.311 Will LPIPS (↓)	geNet PSNR (↑) 17.450 25.226 24.303 23.804 24.278 22.454 xiArt PSNR (↑)	SSIM (†) 0.626 0.800 0.798 0.787 0.777 0.748 SSIM (†)			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired SIGMA Teams VIP	FID (↓) 50.898 10.007 11.770 18.854 12.059 19.646 FID (↓) 75.645	Imag           LPIPS (↓)           0.403           0.207           0.227           0.249           0.217           0.311           Will           LPIPS (↓)           0.437	geNet PSNR (↑) 17.450 25.226 24.303 23.804 24.278 22.454 xiArt PSNR (↑) 17.243	SSIM (†) 0.626 0.800 0.798 0.787 0.777 0.748 SSIM (†) 0.609			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired SIGMA Teams VIP AIIA (Ours)	FID (↓) 50.898 10.007 11.770 18.854 12.059 19.646 FID (↓) 75.645 14.974	Imag           LPIPS (↓)           0.403           0.207           0.227           0.249           0.217           0.311           Will           LPIPS (↓)           0.437           0.244	geNet PSNR (↑) 17.450 25.226 24.303 23.804 24.278 22.454 ciArt PSNR (↑) 17.243 24.350	SSIM (†) 0.626 0.800 0.798 0.787 0.777 0.748 SSIM (†) 0.609 0.767			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired SIGMA Teams VIP AIIA (Ours) HSSLAB	FID (↓) 50.898 10.007 11.770 18.854 12.059 19.646 FID (↓) 75.645 14.974 14.986	Imag           LPIPS (↓)           0.403           0.207           0.227           0.249           0.217           0.311           Will           LPIPS (↓)           0.437           0.244           0.254	geNet PSNR (↑) 17.450 25.226 24.303 23.804 24.278 22.454 ciArt PSNR (↑) 17.243 24.350 24.257	SSIM (↑)           0.626           0.800           0.798           0.787           0.777           0.748           SSIM (↑)           0.609           0.752			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired SIGMA Teams VIP AIIA (Ours) HSSLAB KwaiInpainting	FID (↓) 50.898 10.007 11.770 18.854 12.059 19.646 FID (↓) 75.645 14.974 14.986 26.395	Imag           LPIPS (↓)           0.403           0.207           0.227           0.249           0.217           0.311           Will           LPIPS (↓)           0.437           0.244           0.254           0.276	geNet PSNR (↑) 17.450 25.226 24.303 23.804 24.278 22.454 ciArt PSNR (↑) 17.243 24.350 24.257 23.142	SSIM (↑)           0.626           0.800           0.798           0.787           0.777           0.748           SSIM (↑)           0.609           0.752           0.759			
Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired SIGMA Teams VIP AIIA (Ours) HSSLAB KwaiInpainting ArtificiallyInspired	FID (↓) 50.898 10.007 11.770 18.854 12.059 19.646 FID (↓) 75.645 14.974 14.986 26.395 8.524	Imag           LPIPS (↓)           0.403           0.207           0.227           0.249           0.217           0.311           Will           LPIPS (↓)           0.437           0.254           0.276           0.248	geNet PSNR (↑) 17.450 25.226 24.303 23.804 24.278 22.454 xiArt PSNR (↑) 17.243 24.350 24.257 23.142 23.799	SSIM (†) 0.626 0.800 0.798 0.787 0.777 0.748 SSIM (†) 0.609 0.767 0.752 0.759 0.758			

Table 5. Performance comparisons of different methods on the four datasets. Results are provided from the NTIRE 2022 Image Inpainting challenge report [27].

second stage, we use bilinear to upsample the image and then use another LaMa-like model to refine the upsampled image. We can see from Fig. 1 that the two-stage model called GLaMa-2stage can generate more realistic faces.

### 4.6. NTIRE 2022 Image Inpainting Challenge

This work is proposed to participate in the NTIRE 2022 Image Inpainting Challenge Track 1 Unsupervised [27]. The objective of this challenge is to obtain a mask agnostic network solution capable of producing high-quality results for the best perceptual quality with respect to the ground truth. Our results for the NTIRE 2022 challenge [27] are denoted as AIIA (our team's name) to distinguish from the other teams in Tab. 5. It can be seen that our method goes significantly ahead of FID, LPIPS, PSNR and SSIM for Places2 [42], ImageNet [7] and WikiArt [28] datasets.

#### 5. Conclusion

This study explores the mask generation strategy in image inpainting task. Moreover, we introduce a joint spatial and frequency loss to reduce checkerboard effect and ripples. The proposed GLaMa achieves significant performance improvements over LaMa [31] without changing the model architecture. At the same time, the proposed method is more general and other methods can simply adopt the strategies we have explored. Extensive ablation studies are performed to validate the effectiveness of our method, and comparative results on four public datasets demonstrates the state-of-the-art performance of our method.

**Limitations**. There are also some failed examples when the input image are  $512 \times 512$  pixels, as shown in Fig. 7. The results generated by GLaMa-2stage are sometimes blurred and unreasonable. LaMa [31] and GLaMa also generate low-quality and unpleasant results. In the future, we will incorporate the StyleGAN prior [19] to assist the image inpainting task.

Acknowledgments: The research was supported by the National Natural Science Foundation of China (61971165), in part by the Fundamental Research Funds for the Central Universities (FRFCU5710050119), the Natural Science Foundation of Heilongjiang Province (YQ2020F004).

# References

- Coloma Ballester, Marcelo Bertalmío, Vicent Caselles, Guillermo Sapiro, and Joan Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE Transactions* on *Image Processing*, 10 8:1200–11, 2001. 2
- [2] Marcelo Bertalmío, Guillermo Sapiro, Vicent Caselles, and Coloma Ballester. Image inpainting. In Judith R. Brown and Kurt Akeley, editors, *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, *SIGGRAPH 2000, New Orleans, LA, USA, July 23-28, 2000*, pages 417–424. ACM, 2000. 2
- [3] Marcelo Bertalmío, Luminita A. Vese, Guillermo Sapiro, and S. Osher. Simultaneous structure and texture image inpainting. *IEEE Transactions on Image Processing*, 12 8:882–9, 2003. 2
- [4] Mu Cai, Hong Zhang, Huijuan Huang, Qichuan Geng, and Gao Huang. Frequency domain image translation: More photo-realistic, better identity-preserving. *Proceedings of the IEEE/CVF International Conference on Computer Vision* (*ICCV*), pages 13910–13920, 2021. 3
- [5] Lu Chi, Borui Jiang, and Yadong Mu. Fast fourier convolution. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020. 2, 3
- [6] Antonio Criminisi, Patrick Pérez, and Kentaro Toyama. Object removal by exemplar-based inpainting. volume 2, pages II–II, 2003. 2
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition (CVPR), pages 248–255, 2009. 1, 2, 3, 5, 6, 8
- [8] Qiaole Dong, Chenjie Cao, and Yanwei Fu. Incremental transformer structure enhanced image inpainting with mask-ing positional encoding. *ArXiv*, abs/2203.00867, 2022. 3
- [9] Rinon Gal, Dana Cohen, Amit H. Bermano, and Daniel Cohen-Or. Swagan: A style-based wavelet-driven generative model. ACM Transactions on Graphics (TOG), 40:134:1– 134:11, 2021. 3
- [10] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, 2014. 2, 3
- [11] James Hays and Alexei A. Efros. Scene completion using millions of photographs. ACM Transactions on Graphics (TOG), 2007. 2
- [12] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. pages 770–778, 2016. 4, 6
- [13] 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. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6626–6637, 2017. 6

- [14] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. ACM Transactions on Graphics (TOG), 36:1–14, 2017. 2
- [15] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision (ICCV), pages 13899–13909, 2021. 3, 7
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II, volume 9906 of Lecture Notes in Computer Science, pages 694–711. Springer, 2016. 3
- [17] Steffen Jung and Margret Keuper. Spectral distribution aware image generation. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 1734–1742. AAAI Press, 2021. 3
- [18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *CoRR*, abs/1812.04948, 2018. 1, 2, 3, 4, 5, 6, 7, 8
- [19] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 4401–4410, 2019. 8
- [20] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015. 6
- [21] Guilin Liu, Fitsum A. Reda, Kevin J. Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. *ArXiv*, abs/1804.07723, 2018. 2
- [22] Hongyu Liu, Bin Jiang, Yibing Song, Wei Huang, and Chao Yang. Rethinking image inpainting via a mutual encoderdecoder with feature equalizations. *ArXiv*, abs/2007.06929, 2020. 1, 2
- [23] Kanglin Liu, Wenming Tang, Fei Zhou, and Guoping Qiu. Spectral regularization for combating mode collapse in gans. pages 6381–6389, 2019. 3
- [24] Andreas Lugmayr, Martin Danelljan, Andrés Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. *CoRR*, abs/2201.09865, 2022. 2, 3
- [25] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 5188–5196, 2015. 7

- [26] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting. pages 2536–2544, 2016. 2
- [27] Andres Romero, Angela Castillo, Jose M Abril-Nova, Radu Timofte, et al. NTIRE 2022 image inpainting challenge: Report. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, 2022. 1, 2, 8
- [28] Babak Saleh and A. Elgammal. Large-scale classification of fine-art paintings: Learning the right metric on the right feature. *ArXiv*, abs/1505.00855, 2015. 1, 2, 3, 5, 6, 8
- [29] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *ArXiv*, abs/1503.03585, 2015. 2, 3, 5
- [30] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *ArXiv*, abs/2011.13456, 2021. 2, 3
- [31] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3172–3182, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [32] Ziyu Wan, Bo Zhang, Dongdong Chen, Pan Zhang, Dong Chen, Jing Liao, and Fang Wen. Bringing old photos back to life. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 2747– 2757, 2020. 1
- [33] Ziyu Wan, Jingbo Zhang, Dongdong Chen, and Jing Liao. High-fidelity pluralistic image completion with transformers. arXiv preprint arXiv:2103.14031, 2021. 1, 2
- [34] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 1, 2, 6
- [35] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4470–4479, 2019. 1, 2, 4
- [36] Jiahui Yu, Zhe L. Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S. Huang. Generative image inpainting with contextual attention. pages 5505–5514, 2018. 2
- [37] 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. In *Proceedings of the* 29th ACM International Conference on Multimedia, 2021. 1, 2
- [38] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Baining Guo. Aggregated contextual transformations for highresolution image inpainting. *IEEE Transactions on Visualization and Computer Graphics*, PP, 2022. 1, 2
- [39] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. 1, 2, 6

- [40] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I Chang, and Yan Xu. Large scale image completion via co-modulated generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2021. 1
- [41] Shengyu Zhao, Jianwei Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I-Chao Chang, and Yan Xu. Large scale image completion via co-modulated generative adversarial networks. *ArXiv*, abs/2103.10428, 2021. 2
- [42] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analy*sis and Machine Intelligence, 2017. 1, 2, 3, 5, 6, 8
- [43] Manyu Zhu, Dongliang He, Xin Li, Chao Li, Fu Li, Xiao Liu, Errui Ding, and Zhaoxiang Zhang. Image inpainting by end-to-end cascaded refinement with mask awareness. *IEEE Transactions on Image Processing*, 30:4855–4866, 2021. 2