

CoordFill: Efficient High-Resolution Image Inpainting via Parameterized Coordinate Querying

Wei Huang Liu¹, Xiaodong Cun^{2*}, Chi-Man Pun^{1*}, Menghan Xia², Yong Zhang², Jue Wang²

¹ University of Macau,

² Tencent AI Lab

mc05379@umac.mo, vinthony@gmail.com, cmpun@umac.mo,
menghanxyz@gmail.com, zhangyong201303@gmail.com, arphid@gmail.com

Abstract

Image inpainting aims to fill the missing hole of the input. It is hard to solve this task efficiently when facing high-resolution images due to two reasons: (1) Large reception field needs to be handled for high-resolution image inpainting. (2) The general encoder and decoder network synthesizes many background pixels synchronously due to the form of the image matrix. In this paper, we try to break the above limitations for the first time thanks to the recent development of continuous implicit representation. In detail, we down-sample and encode the degraded image to produce the spatial-adaptive parameters for each spatial patch via an attentional Fast Fourier Convolution (FFC)-based parameter generation network. Then, we take these parameters as the weights and biases of a series of multi-layer perceptron (MLP), where the input is the encoded continuous coordinates and the output is the synthesized color value. Thanks to the proposed structure, we only encode the high-resolution image in a relatively low resolution for larger reception field capturing. Then, the continuous position encoding will be helpful to synthesize the photo-realistic high-frequency textures by re-sampling the coordinate in a higher resolution. Also, our framework enables us to query the coordinates of missing pixels only in parallel, yielding a more efficient solution than the previous methods. Experiments show that the proposed method achieves real-time performance on the 2048×2048 images using a single GTX 2080 Ti GPU and can handle 4096×4096 images, with much better performance than existing state-of-the-art methods visually and numerically. The code is available at: <https://github.com/NiFangBaAGe/CoordFill>.

Introduction

With the rapid development of digital media, the demand for image editing, for example, filling the given hole according to the remaining background pixels or removing the unnecessary object, also increases a lot. These tasks belong to image inpainting, which is one of the fundamental image synthesis tasks, which requires both semantic understanding and conditional generation.

Convolutional neural network (CNN)-based methods dominate this field in recent years. For example, coarse-to-fine network structures and generative adversarial networks (Goodfellow et al. 2014) based methods (Yu et al.

*Corresponding Author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

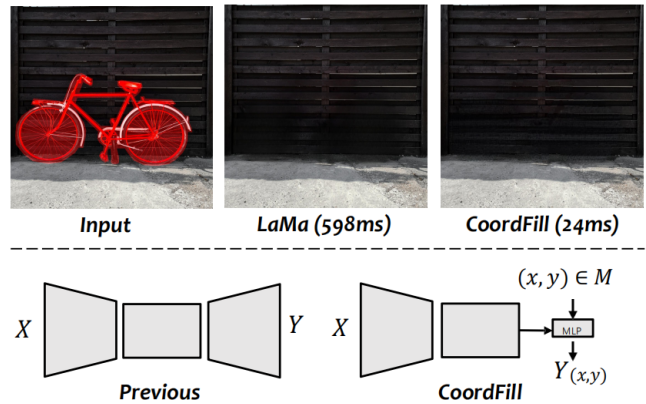


Figure 1: The proposed method only requires 24ms ($25 \times$ faster than state-of-the-art method LaMa (Suvorov et al. 2022) in the same resolution) to remove the red masked regions on the resolution of a 2048×2048 image. Compared with most of the previous methods that synthesize the full image Y by convolutional neural networks, CoordFill queries the coordinates (x, y) in the missing mask M of the input X only and generates the pixel-value by the implicit representation. Thanks to the proposed decoder, CoordFill runs faster than previous methods on the high-resolution image inpainting task.

2018, 2019; Iizuka, Simo-Serra, and Ishikawa 2017; Liu et al. 2018) are used to synthesize the realistic textures. Other types of image inpainting are inspired by the typical hints (e.g. edges (Nazeri et al. 2019), semantic maps (Liao et al. 2021)) from the original background, these methods guess the hints in the hole region firstly, and then, they synthesize the texture by the guidance of the filled hints. However, current methods still suffer from many limitations for real-world applications since casual images are in high resolution and on the mobile platform. We raise a question: *What makes the inpainting hard in efficiently handling the high-resolution images?* and we attempt to answer this question from the following perspectives:

(i) Learning the larger reception fields is time-consuming for high-resolution images. A larger reception field is essential in image inpainting (Suvorov et al. 2022), which means that we need to feed the full-size images to the inpainting

network. In general, there are two straightforward ways to increase speed. We can crop or resize the images to fit the resolution of the input. However, direct cropping will influence the semantic guidance from the background region and hugely reduce the ability of scene understanding. Directly resizing is another option, it requires preserving the detailed high-frequency details for upsampling.

(ii) The decoders will also synthesize a lot of unnecessary pixels. Image inpainting aims to remove a relatively small hole from the input. Restricted by the form of the image matrix, small irregular holes also need to be synthesized using a stack of convolutional layers in full image size.

We try to break the above limitations via the implicit representation from neural rendering community (Sitzmann et al. 2020; Mildenhall et al. 2020; Shaham et al. 2021). A typical implicit representation overfits the scene to small multi-layer perceptrons (MLPs) network, where the input is the coordinate of the image matrix and the output is the queried pixels. It is a continuous representation that can keep the high-frequency details well and can achieve pixel-wise reconstruction. Inspired by these wonderful features, we propose a novel framework for efficient high-resolution image inpainting using implicit representation for the first time (Figure 1). To make the network related to the specific input, our method utilize a meta-learning strategy where the parameters of the pixel-wise query network are spatial-adaptive (Shaham et al. 2021). In detail, our framework contains two sub-networks for parameter generation and pixel-wise query. In the parameter generation network, we resize the high-resolution image and generate the parameters for each local image patch. This network is built via a series of the proposed attentional FFC blocks. Then, the pixel-wise query network reforms the predicted parameters as a series of the MLP, where the input is the positional embeddings in the frequency domain and the output is the predicted color value in the given coordinate. Thanks to the proposed structure, our method can run faster in the challenge of high-resolution image inpainting from two aspects. On the one hand, most of the costly operations are in the low-resolution parameter generation network, where we can learn larger reception fields. On the other hand, we can decode the masked regions only since the coordinate can be queried one-by-one in parallel and the larger size image can be got via the coordinate re-sampling. The detailed experiments show the efficiency and the power of the proposed method compared with previous state-of-the-art on several benchmarks.

The contributions of this paper are summarized as follows:

- We propose CoordFill, a novel framework for efficient high-resolution image inpainting via parameterized coordinate querying.
- We design an attentional FFC-based block as the basic block in our parameter generation network. It learns to focus on the masked region automatically.
- Our method runs faster than previous baselines and achieves state-of-the-art performance in image inpainting on multiple datasets.

Related Work

Image Inpainting Traditional image inpainting methods rely on strong low-level assumptions, for example, local patches (Barnes et al. 2009) can be used to fill the missing region. Several recent CNN-based methods (Liu et al. 2020; Suvorov et al. 2022; Iizuka, Simo-Serra, and Ishikawa 2017; Yu et al. 2018; Zeng et al. 2021; Yu et al. 2019; Wan et al. 2021; Liu et al. 2018; Yu et al. 2020; Zhang et al. 2018a; Vo, Duong, and Pérez 2018; Li et al. 2022; Dong, Cao, and Fu 2022) have a similar or a stack of encoder-decoder architectures. More specifically, Iizuka *et al.* (Iizuka, Simo-Serra, and Ishikawa 2017) introduce a GAN-based framework by the global and local discriminator. Based on the GAN framework, more novel blocks are also introduced through Attention (Yu et al. 2018, 2019; Liu et al. 2020; Zeng et al. 2021), Regional Normalization (Yu et al. 2020) and Regional Convolutions (Liu et al. 2018). Novel structures, like Vision Transformer (Dosovitskiy et al. 2020), also draw the attention of the image editing community (Wan et al. 2021). Meaningful priors, such as edge (Nazeri et al. 2019) and semantic label (Liao et al. 2021) also play an important role. For high-resolution image inpainting, a multi-stage network is a common choice. For example, Yi *et al.* (Yi et al. 2020) refine the low-resolution output via the contextual residual aggregation. Zeng *et al.* (Zeng et al. 2020b) design a multi-scale network structure with guided up-sampling. MAT (Li et al. 2022) and ZITS (Dong, Cao, and Fu 2022) are transformer-based image inpainting systems for high-resolution images. However, the multi-stage network structure is slow for real-world applications. In contrast, LaMa (Suvorov et al. 2022) designs a one-stage network by the fusion of the multi-scale reception fields. However, their speed is also restricted by the normal decoder.

Continuous Image Representation Learning the continuous image representation, *i.e.*, the implicit representation, are popular recently, due to the success in 3D human digitization (Saito et al. 2019, 2020) and novel view synthesis (Mildenhall et al. 2020). Generally, these methods query the color of the specific localization using the positional encoding in the frequency domain, which can generate more detailed high-frequency details than other methods. However, implicit representation is rarely developed in the 2D image domain. StyleGANv3 (Karras et al. 2021) is built on implicit representation for alias-free settings of the image generation in the specific domain. InfinityGAN (Lin et al. 2021) also learns for image synthesis. ASAPNet (Shaham et al. 2021) learns the pixels from the semantic layout. LIIF (Chen, Liu, and Wang 2021) learns a continuous representation for image super-resolution. However, how to utilize this technique in image inpainting is still unclear and we make the first step.

Efficient Network Structure Designing an efficient network structure for the high-resolution image is a more practical consideration for any deep learning-based method. Several attempts have also been made to make the convolutional network work in mobile applications. For example, the efficient network backbone has been widely explored in image classification (Howard et al. 2017; Iandola et al. 2016). In

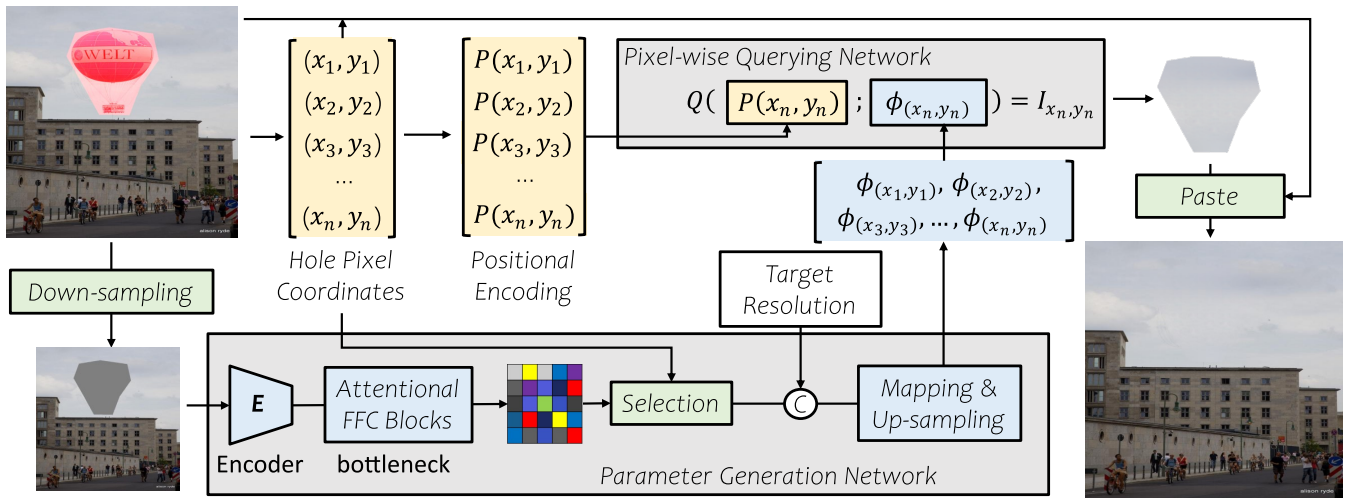


Figure 2: Overview of the proposed framework.

the image editing and synthesis fields, the models can be distilled from a trained larger network using knowledge distillation (Li et al. 2020). For the image retouching tasks (Zeng et al. 2020a; Liang, Cun, and Pun 2021), learning the global-aware representation is helpful to keep the efficiency in high-resolution images. As for the high-resolution inpainting, efficiency is not the concern of the community yet.

Method

Overview Our goal is to synthesize the missing content for high-resolution images in an efficient manner and hold a stable performance to diverse resolutions. To this end, our proposed model, dubbed as CoordFill, solves the inpainting problem with two phases: (i) restorative feature synthesis, as represented by spatially-adaptive parameters of the reconstruction function; (ii) individual pixel reconstruction based on the per-pixel coordinate query. The advantage of this design is that we can synthesize high-resolution content efficiently for any specified regions thanks to the coordinate query-based reconstruction scheme. Therefore, on the one hand, it is allowed to selectively reconstruct the hole region of the input image. On the other hand, we can resample the input image with a fixed resolution before being fed to the network, which enables the contextual feature extraction under a unified receptive field, yet the image still could be restored with the original resolution (i.e. by sub-pixel coordinate query).

Specifically, given a high-resolution masked image $\mathbf{I} * \mathbf{M} \in \mathbb{R}^{H \times W \times 3}$ and its mask $\mathbf{M} \in [0, 1]^{H \times W}$, we first downsample them by S times to 256×256 . Then, we build a Parameter Generation Network G to generate the full image representation as a parameter map Φ . Finally, a Pixel-wise Querying Network Q , parameterized by Φ_p , generates pixel values according to the coordinate query \mathbf{p} at the corresponding patches and finally generates the restored image \mathbf{I}_o by pasting back the non-hole content. The overall architecture is illustrated in Figure 2. Below, we give the details of each component.

Parameter Generation Network Existing inpainting models are typically formulated to generate a full image with hole regions filled but only the hole region pixels are used in the final results. Obviously, the non-hole pixels are not necessary to be synthesized, as they contribute nothing to the final visualization. To avoid such unnecessary computation, we propose to first generate a parameter map from the input and then reconstruct the pixel values at hole regions individually. As the inpainting task relies more on the surrounding context interpretation than texture details, we propose to resample the input with a fixed resolution before feeding it to our parameter generation network G . This operation not only promotes computation efficiency, especially for high-resolution images but also guarantees our network with good generalization to various input resolutions. Note that, we can still reconstruct the hole region with the original resolution, due to the coordinate query-based reconstruction scheme.

We employ a fully convolutional structure for the parameter generation network, composed of the encoder, bottleneck blocks, and the final layers of our parameter generation network. In detail, we down-sample the input feature via the convolution layer three times to favor the process of calculation.

In bottleneck blocks, inspired by recently popular Fast Fourier Convolution (FFC) (Chi, Jiang, and Mu 2020; Suvorov et al. 2022) blocks, we propose an FFC-based framework also. This block captures the global reception fields of the input in the frequency domain, yet a better performance in image inpainting which needs larger reception fields. Please refer to the previous works (Chi, Jiang, and Mu 2020; Suvorov et al. 2022) for more details.

However, the global receptive field in FFC is only limited in the frequency domain, which lacks representation in the spatial space. To alleviate this problem, we propose a spatial-aware attention block based on FFC (AttFFC). The proposed block contains $-/+$ operations to remove the spatial noise and enhance the features (Zheng et al. 2019). Given a feature map $\mathbf{F} \in \mathbb{R}^{H \times W \times C}$, AttFFC first computes the spa-

tial attention map \mathbf{F}_{dm} through a FFC Block and a Sigmoid function σ :

$$\mathbf{F}_{dm} = \sigma(\text{FFC}(\mathbf{F})). \quad (1)$$

where F_{dm} are the one channel spatial confidential map automatically. Next, we use this spatial prior and the original feature to get the spatial noise:

$$\mathbf{F}_d = \text{FFC}(\text{concat}(\mathbf{F}, \mathbf{F}_{dm})). \quad (2)$$

Then, we remove the spatial noise, and enhance the features by learning the residual with an FFC block:

$$\mathbf{F}' = \mathbf{F} - \mathbf{F}_d + \text{FFC}(\mathbf{F} - \mathbf{F}_d). \quad (3)$$

We stack six AttFFCs in our parameter generation network as the bottleneck.

Notice that, since we use the encoder to reduce the feature space, the learned spatial feature is also in a small resolution. To this end, we build the pixel-wise parameters via selection, mapping, and upsampling. In detail, we first select the masked region features by the corresponding coordinates, and then, for each spatial feature, we generate the required number of the parameters by the linear mapping function f . In addition, to make the parameter generation network sensitive to the target resolution r , we inject the target resolution as a conditional input to f :

$$\phi_m = f(\text{cat}[F_m, r]), \quad (4)$$

where F_m is the features in the masked region only, and $\text{cat}[\dots]$ denotes the concatenation operation. Finally, the spatial-adaptive patch ϕ_m will be upsampled to the required resolution with nearest-neighbor interpolation. Hence, we obtain a series of parameter vectors with channels corresponding to the weights (and the biases) of the MLPs, and the heavy computation is performed at low resolution.

Pixel-wise Querying Network As mentioned above, we adopt a coordinate query-based pixel reconstruction scheme by utilizing the spatially-adaptive parameters. In particular, inspired by the efficiency of implicit representation (Shaham et al. 2021), we employ a Multi-Layer Perception (MLP) for per-pixel reconstruction, *i.e.* the Pixel-wise Querying Network Q , which is parameterized by the spatial-adaptive contextual information embedded parameters, as extracted by G . Note that, Q is just a function form and its expressive capability mainly inherits from those deeply extracted parameters Φ . Formally, Q takes the pixel’s coordinate \mathbf{p} as input and outputs the color values:

$$\mathbf{y}_p = Q(\mathbf{p}; \Phi_p), \quad (5)$$

To boost high-frequency details, we encode each component of the 2D pixel position $p = (P_x, P_y)$ as a vector of sinusoids (Vaswani et al. 2017). Additionally, we normalize the positional encoding system and sample positional encoding within a fixed range. Thus, the pixels in the high-resolution images can be synthesized by changing the interval of the positional encoding. For example, given an input image \mathbf{I} and its hole mask \mathbf{M} of resolution $H \times W$, our parameter generation network generates the spatially-adaptive parameter map Φ of resolution $h \times w$ from the downsampled input

$\{\tilde{\mathbf{I}}, \tilde{\mathbf{M}}\}$. Then, we feed the coordinate query:

$$\mathbf{p} = (\sin(2\pi p_x/\mathbf{E}_x), \cos(2\pi p_x/\mathbf{E}_x), \sin(2\pi p_y/\mathbf{E}_y), \cos(2\pi p_y/\mathbf{E}_y)) \quad (6)$$

at expected interval $\mathbf{E}_x, \mathbf{E}_y$, where $\mathbf{E}_x = \frac{H}{h}$ and $\mathbf{E}_y = \frac{W}{w}$, and then synthesize the hole-region content of the original resolution.

Loss Function Since image inpainting only needs to measure the realism of the content, we train our model with different perception losses, including the differences in the pre-trained AlexNet domain (Zhang et al. 2018b; Jo, Yang, and Kim 2020), adversarial loss (Goodfellow et al. 2014), and feature matching loss (Suvorov et al. 2022; Wang et al. 2018).

For the perceptual loss on the pre-trained ImageNet classification domain, the loss function is defined as:

$$L_{per} = \sum_k \tau_k (E_k(\mathbf{I}_o) - E_k(\mathbf{I}_{gt})), \quad (7)$$

where E is an AlexNet feature extractor, τ calculates the differences in the feature domain via L_1 loss, and then, we compute and average the losses from k layers.

Then, the adversarial loss is used to encourage the model to generate more realistic details, which can be written as:

$$L_{adv} = \mathbb{E}[\log(1 - D(\mathbf{I}_o))] + \mathbb{E}[\log D(\mathbf{I}_{gt})]. \quad (8)$$

where \mathbb{E} denotes expectation values over the training batch.

Next, the feature matching loss is adopted for stabilizing the GAN training (Wang et al. 2018; Suvorov et al. 2022):

$$L_{fm} = \sum_i (D^i(\mathbf{I}_{gt}), D^i(\mathbf{I}_o)), \quad (9)$$

where D^i denotes the activations from the i -th layer of the discriminator D .

Finally, the total loss of our model can be written as:

$$L_{total} = \lambda_{per} L_{per} + \lambda_{adv} L_{adv} + \lambda_{fm} L_{fm}. \quad (10)$$

where we empirically set $\lambda_{per} = 10$, $\lambda_{adv} = 1$ and $\lambda_{fm} = 100$ respectively.

Experiments

Implementation Details

We evaluate our method on the widely-used Places2 (Zhou et al. 2017) dataset, as well as two high-resolution image datasets, CelebA-HQ (Karras et al. 2018) and Unsplash (Unsplash 2021). We use the irregular mask from the previous method (Liu et al. 2018) for mask generation with holes up to 25% following HiFill (Yi et al. 2020). Our model is optimized by the Adam optimizer with a learning rate of $1e-4$ and trained with 100 epochs. All the training experiments are conducted on 8 Tesla A100 GPUs with a batch size of 128. All the images are randomly resized from 256×256 to 512×512 during training, then, we test them in the range of 512, 1024, 2048, and 4096. Following previous methods (Yi et al. 2020; Zeng et al. 2022), we use the PSNR, SSIM, and LPIPS (Zhang et al. 2018b) for performance comparison.



Figure 3: Qualitative comparisons of the proposed CoordFill and other methods. The top two examples are 512×512 images from the Places2 dataset while the bottom two images are 1024×1024 images from the Unsplash dataset.

Resolution	512×512			1024×1024			2048×2048			4096×4096		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
DeepFillv2	23.973	0.902	0.080	22.695	0.908	0.092	-	-	-	-	-	-
HiFill	23.375	0.883	0.097	23.456	0.894	0.096	23.643	0.915	0.087	23.634	0.933	0.077
RN	22.562	0.880	0.116	19.587	0.879	0.139	18.843	0.908	0.143	-	-	-
CR-Fill	24.216	0.893	0.086	22.881	0.890	0.108	22.056	0.908	0.122	-	-	-
LaMa	26.203	0.914	0.067	26.154	0.924	0.076	25.688	0.939	0.078	-	-	-
MAT	24.169	0.900	0.076	23.751	0.908	0.082	-	-	-	-	-	-
ZITS	26.349	0.911	0.068	26.389	0.913	0.073	-	-	-	-	-	-
CoordFill	26.365	<u>0.912</u>	<u>0.068</u>	<u>26.322</u>	<u>0.920</u>	<u>0.075</u>	26.322	<u>0.932</u>	0.077	26.175	0.943	0.075

Table 1: Comparison with state-of-the-art methods on Places2 Dataset on different resolutions. “-” means these methods cause the out-of-memory issue in the given resolutions.

Comparisons with Prior Arts

Our method is compared against the existing state-of-the-art image inpainting algorithms including the general image inpainting models (DeepFillv2 (Yu et al. 2019), RN (Yu et al. 2020) and CR-Fill (Zeng et al. 2021)), and the high-resolution image inpainting models (HiFill (Yi et al. 2020), LaMa (Suvorov et al. 2022), MAT (Li et al. 2022), and ZITS (Dong, Cao, and Fu 2022)) using their official pre-trained models. As shown in Table 1, our CoordFill achieves competitive performance at all resolutions on the Places2 validation set. Compared with the common image inpainting methods (DeepFillv2, RN, and CR-Fill) on the 512×512 images, the proposed method achieves the best performance. Compared with the existing high-resolution image inpainting methods (HiFill, LaMa, MAT, and ZITS) on various resolutions, the proposed method also shows more stable results when the resolution increases thanks to the robust resolution and continuous pixel-wise query. For visual comparison, we also compare our methods with other state-of-the-art methods in Figure 3. According to the figure, the proposed method shows more visual-please results than the previous

methods with stable reception fields and pixel-wise querying.

Since the previous high-resolution image inpainting method (Yi et al. 2020) generates the high-resolution samples (>512) by the bilinear upsampling on Places2, we use another two high-resolution image datasets to evaluate the performance of each method more accurately, including CelebA-HQ and Unsplash. CelebA-HQ is a commonly used high-resolution face dataset with a resolution of 1024×1024 , and the model is retrained with its train set and evaluated on its test set. Unsplash is a real-world high-resolution dataset collected from Unsplash (Unsplash 2021) in which the original images are cropped to 1024×1024 , and we use the pre-trained model on Place2 for evaluation. Table 2 reports the performance of each method on these datasets, where CoordFill shows competitive performance. As for the visual comparison, the proposed CoordFill also generates clear details as in Figure 3.

	Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Unsplash	DeepFillv2	24.859	0.914	0.097
	HiFill	27.653	0.914	0.082
	RN	20.290	0.890	0.162
	CR-Fill	28.690	0.923	0.082
	LaMa	32.434	0.934	0.055
	MAT	29.506	0.923	0.081
	ZITS	32.532	0.934	0.053
	CoordFill	33.093	0.934	0.080
CelebA-HQ	DeepFillv2	23.286	0.916	0.106
	LaMa	26.098	0.924	0.080
	MAT	25.167	0.917	0.074
	CoordFill	28.756	0.934	0.065

Table 2: Comparison with state-of-the-art methods on Unsplash and CelebA-HQ Datasets (resolution 1024×1024).

	512×512	1024×1024	2048×2048	4096×4096
DeepFillv2	398	1002	-	-
HiFill	406	423	478	662
RN	17	59	249	-
CR-Fill	46	54	63	-
LaMa	27	142	598	-
MAT	71	133	-	-
ZITS	183	462	-	-
CoordFill	10	14	26	78

Table 3: Speed comparison (ms) with state-of-the-art methods on different resolutions. “-” means these methods cause the out-of-memory issue in the given resolutions.

Masked Prediction	Decoder	Block	Resolution Injection	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
-	D_{Conv}	ResFFC	-	25.5420	0.9074	0.0688
✓	D_{Conv}	ResFFC	-	25.9557	0.9097	0.0679
✓	D_{MLP}	ResFFC	-	24.6480	0.8995	0.1186
-	PQN	ResFFC	-	26.0434	0.9091	0.0683
✓	PQN	ResFFC	-	26.2160	0.9109	0.0679
✓	PQN	AttFFC	-	26.3177	0.9116	0.0679
✓	PQN	AttFFC	✓	26.3653	0.9119	0.0677

Table 4: Ablation study on the components in the proposed framework. The experiments are conducted on the 512×512 Place2 images.

Efficiency Analysis

We give a detailed efficiency analysis. Firstly, we show the speed comparison of our method and others in Table 3. The speed is calculated using the average inference time on an NVIDIA GTX 2080 Ti GPU. Specifically, the proposed method runs much faster than all the methods under all the resolutions. On the same experiment platform, nearly all the methods cause out-of-memory issues when the resolution increases to 4096×4096 except ours and Hi-Fill. Also, as shown in Table 1, the proposed method shows a much better performance.

Furthermore, as shown in Figure 4, we give a detailed

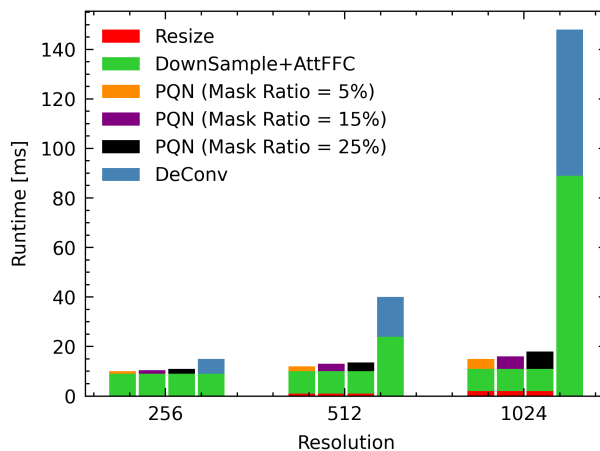


Figure 4: The speed comparison of the proposed method (three different mask ratios) and the baseline (DownSample + AttFFC + D_{Conv}) on the three different resolutions.




Figure 5: Qualitative comparisons of different supervision strategies.

explanation of what exactly influences the speed when the resolution changes. As in Figure 4, our pixel-wise querying network always runs faster than the normal convolutional decoder (D_{Conv}). On the other hand, the speedup comes from the stable reception fields where larger feature maps will speed more time. It is also interesting to see that the running time of our method is flexible when the hole changes. It is because our method only needs to predict the masked region using the pixel-query strategy.




Figure 6: Qualitative comparisons of decoding using a shared MLP, CNN, and the pixel-wise querying network (PQN).

Ablation Study

We conduct a detailed ablation study to verify the efficiency of each component in the proposed framework on the 512×512 Places2 images, and the quantitative results are shown in Table 4.

Supervision strategy We design a novel framework for image inpainting, *i.e.* selective region synthesis based on the coordinate query. Interestingly, we find that existing methods still calculate the loss function on the whole image, which might be biased supervision because we only care about the hole region accuracy. To verify this speculation, we calculate the loss function in the masked region only, called *Masked Prediction*. As shown in Table 4, calculating the loss function in the masked region improves the performance in both D_{Conv} and the proposed pixel-wise querying network (PQN) structure on both two datasets. It might be because the original supervision strategy for the full image synthesis will be influenced by the extra reception fields and result in color bleeding artifacts as shown in Figure 5.

Importance of the coordinate-query scheme The proposed pixel-wise querying network helps our method to obtain high-resolution and high-quality results, which is the key to our method. Thus, we compare it with D_{Conv} and D_{MLP} with the same encoder, where D_{Conv} uses the original convolutional decoder for up-sampling and D_{MLP} feed the pixel-wise feature to a shared MLP for pixel value decoding. As shown in Table 4 and Figure 6, the proposed pixel-wise querying framework gets much better performance and runs faster than the previous, which proves it is an efficient approach for high-resolution image inpainting.

Advantage of our attentional FFC The proposed attentional FFC (AttFFC) is also an important component. As we have claimed, in AttFFC, we suppress the spatial noise by filtering out the distracting features. Compared with the orig-




Figure 7: Qualitative comparisons of CoordFill and super-resolution methods on real-world high-resolution images.

inal FFC blocks (ResFFC) in LaMa (Suvorov et al. 2022), AttFFC achieves a better result than ResFFC as shown in Table 4.

Impact of resolution injection To extract features under a unified receptive field, the parameter generation network accepts a fixed low-resolution input. Hence the pixel-wise querying network generates the same results for different resolutions, and the high-frequency information can only be synthesized with the help of the input positional encoding. By condition on the resolution information, our method got a better performance as shown in Table 4.

Comparisons with inpainting then super-resolution

Our coordinate query-based decoding network supports synthesizing the content in arbitrary resolution, and it is also implicitly required to consider the super-resolution as post-processing during testing. To check the effectiveness, we construct a baseline: the decoding network only generates the output of the same resolution as the down-sampled input (256×256) and then gets high-resolution image (1024×1024 in our case) by an extra up-sampling or super-resolution network. We plot the visual results to support our claim in Figure 7, where the proposed method shows sharper and visual-friendly results.

Conclusion

We use the continuous coordinate representation for efficient high-resolution image inpainting. To this end, we propose CoordFill, a novel framework that contains the parameter generation network and the pixel-wise querying network following the meta-learning strategy. In detail, in the parameter generation network, we design a novel attention network based on FFC to produce the parameters in a spatial-adaptive fashion. In the pixel-wise querying network, we use the generated parameters from the parameter generation network and query the masked coordinate by re-sampling the coordinates in the frequency domain. The experiments show the efficiency of the proposed method on several datasets compared with several state-of-the-art methods.

Acknowledgments

This work was supported in part by the Science and Technology Development Fund, Macau SAR, under Grant 0034/2019/AMJ, Grant 0087/2020/A2, and Grant 0049/2021/A.

References

- Barnes, C.; Shechtman, E.; Finkelstein, A.; and Goldman, D. B. 2009. PatchMatch: A Randomized Correspondence Algorithm for Structural Image Editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 28(3).
- Chen, Y.; Liu, S.; and Wang, X. 2021. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8628–8638.
- Chi, L.; Jiang, B.; and Mu, Y. 2020. Fast fourier convolution. *Advances in Neural Information Processing Systems*, 33: 4479–4488.
- Dong, Q.; Cao, C.; and Fu, Y. 2022. Incremental Transformer Structure Enhanced Image Inpainting with Masking Positional Encoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
- Iandola, F. N.; Han, S.; Moskewicz, M. W.; Ashraf, K.; Dally, W. J.; and Keutzer, K. 2016. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and 0.5 MB model size. *arXiv preprint arXiv:1602.07360*.
- Iizuka, S.; Simo-Serra, E.; and Ishikawa, H. 2017. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4): 1–14.
- Jo, Y.; Yang, S.; and Kim, S. J. 2020. Investigating loss functions for extreme super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 424–425.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *International Conference on Learning Representations*.
- Karras, T.; Aittala, M.; Laine, S.; Härkönen, E.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2021. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34.
- Li, M.; Lin, J.; Ding, Y.; Liu, Z.; Zhu, J.-Y.; and Han, S. 2020. GAN Compression: Efficient Architectures for Interactive Conditional GANs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Li, W.; Lin, Z.; Zhou, K.; Qi, L.; Wang, Y.; and Jia, J. 2022. MAT: Mask-Aware Transformer for Large Hole Image Inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Liang, J.; Cun, X.; and Pun, C.-M. 2021. Spatial-Separated Curve Rendering Network for Efficient and High-Resolution Image Harmonization. *arXiv preprint arXiv:2109.05750*.
- Liao, L.; Xiao, J.; Wang, Z.; Lin, C.-W.; and Satoh, S. 2021. Image inpainting guided by coherence priors of semantics and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6539–6548.
- Lin, C. H.; Lee, H.-Y.; Cheng, Y.-C.; Tulyakov, S.; and Yang, M.-H. 2021. InfinityGAN: Towards Infinite-Pixel Image Synthesis. *arXiv preprint arXiv:2104.03963*.
- Liu, G.; Reda, F. A.; Shih, K. J.; Wang, T.-C.; Tao, A.; and Catanzaro, B. 2018. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European conference on computer vision (ECCV)*, 85–100.
- Liu, H.; Jiang, B.; Song, Y.; Huang, W.; and Yang, C. 2020. Rethinking image inpainting via a mutual encoder-decoder with feature equalizations. In *European Conference on Computer Vision*, 725–741. Springer.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, 405–421. Springer.
- Nazeri, K.; Ng, E.; Joseph, T.; Qureshi, F. Z.; and Ebrahimi, M. 2019. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212*.
- Saito, S.; Huang, Z.; Natsume, R.; Morishima, S.; Kanazawa, A.; and Li, H. 2019. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2304–2314.
- Saito, S.; Simon, T.; Saragih, J.; and Joo, H. 2020. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 84–93.
- Shaham, T. R.; Gharbi, M.; Zhang, R.; Shechtman, E.; and Michaeli, T. 2021. Spatially-adaptive pixelwise networks for fast image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14882–14891.
- Sitzmann, V.; Martel, J.; Bergman, A.; Lindell, D.; and Wetzstein, G. 2020. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33: 7462–7473.
- Suvorov, R.; Logacheva, E.; Mashikhin, A.; Remizova, A.; Ashukha, A.; Silvestrov, A.; Kong, N.; Goka, H.; Park, K.; and Lempitsky, V. 2022. Resolution-robust Large Mask Inpainting with Fourier Convolutions. In *Proceedings of the*

- IEEE/CVF Winter Conference on Applications of Computer Vision*, 2149–2159.
- Unsplash. 2021. Unsplash Lite Dataset 1.2.0. <https://unsplash.com/data>. Accessed: 2021-03-04.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Vo, H. V.; Duong, N. Q.; and Pérez, P. 2018. Structural inpainting. In *Proceedings of the 26th ACM international conference on Multimedia*, 1948–1956.
- Wan, Z.; Zhang, J.; Chen, D.; and Liao, J. 2021. High-Fidelity Pluralistic Image Completion With Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 4692–4701.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Yi, Z.; Tang, Q.; Azizi, S.; Jang, D.; and Xu, Z. 2020. Contextual residual aggregation for ultra high-resolution image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7508–7517.
- Yu, J.; Lin, Z.; Yang, J.; Shen, X.; Lu, X.; and Huang, T. S. 2018. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5505–5514.
- Yu, J.; Lin, Z.; Yang, J.; Shen, X.; Lu, X.; and Huang, T. S. 2019. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4471–4480.
- Yu, T.; Guo, Z.; Jin, X.; Wu, S.; Chen, Z.; Li, W.; Zhang, Z.; and Liu, S. 2020. Region normalization for image inpainting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12733–12740.
- Zeng, H.; Cai, J.; Li, L.; Cao, Z.; and Zhang, L. 2020a. Learning Image-adaptive 3D Lookup Tables for High Performance Photo Enhancement in Real-time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zeng, Y.; Fu, J.; Chao, H.; and Guo, B. 2022. Aggregated contextual transformations for high-resolution image inpainting. *IEEE Transactions on Visualization and Computer Graphics*.
- Zeng, Y.; Lin, Z.; Lu, H.; and Patel, V. M. 2021. Cr-fill: Generative image inpainting with auxiliary contextual reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14164–14173.
- Zeng, Y.; Lin, Z.; Yang, J.; Zhang, J.; Shechtman, E.; and Lu, H. 2020b. High-resolution image inpainting with iterative confidence feedback and guided upsampling. In *European conference on computer vision*, 1–17. Springer.
- Zhang, H.; Hu, Z.; Luo, C.; Zuo, W.; and Wang, M. 2018a. Semantic image inpainting with progressive generative networks. In *Proceedings of the 26th ACM international conference on Multimedia*, 1939–1947.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018b. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zheng, Q.; Qiao, X.; Cao, Y.; and Lau, R. W. 2019. Distraction-aware shadow detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5167–5176.
- Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6): 1452–1464.