

Contextual Residual Aggregation for Ultra High-Resolution Image Inpainting

Zili Yi Qiang Tang Shekoofeh Azizi Daesik Jang Zhan Xu
Huawei Technologies Canada Co. Ltd.

{zili.yi, qiang.tang, shekoofeh.azizi, daesik.jang, zhan.xu}@huawei.com




Figure 1: Inpainting results on ultra high-resolution images.

Abstract

Recently data-driven image inpainting methods have made inspiring progress, impacting fundamental image editing tasks such as object removal and damaged image repairing. These methods are more effective than classic approaches, however, due to memory limitations they can only handle low-resolution inputs, typically smaller than 1K. Meanwhile, the resolution of photos captured with mobile devices increases up to 8K. Naive up-sampling of the low-resolution inpainted result can merely yield a large yet blurry result. Whereas, adding a high-frequency residual image onto the large blurry image can generate a sharp result, rich in details and textures. Motivated by this, we propose a Contextual Residual Aggregation (CRA) mechanism that can produce high-frequency residuals for missing contents by weighted aggregating residuals from contextual patches, thus only requiring a low-resolution prediction from the network. Since convolutional layers of

the neural network only need to operate on low-resolution inputs and outputs, the cost of memory and computing power is thus well suppressed. Moreover, the need for high-resolution training datasets is alleviated. In our experiments, we train the proposed model on small images with resolutions 512×512 and perform inference on high-resolution images, achieving compelling inpainting quality. Our model can inpaint images as large as 8K with considerable hole sizes, which is intractable with previous learning-based approaches. We further elaborate on the light-weight design of the network architecture, achieving real-time performance on 2K images on a GTX 1080 Ti GPU. Codes are available at: https://github.com/Ascend-Huawei/Ascend-Canada/tree/master/Models/Research_HiFill_Model

1. Introduction

Smartphone users are interested to manipulate their photographs in any form of altering object positions, remov-

ing unwanted visual elements, or repairing damaged images. These tasks require automated image inpainting, which aims at restoring lost or deteriorated parts of an image given a corresponding mask. Inpainting has been an active research area for the past few decades, however, due to its inherent ambiguity and the complexity of natural images, general image inpainting remains challenging. High-quality inpainting usually requires generating visually realistic and semantically coherent content to fill the hole regions. Existing methods for image hole filling can be categorized into three groups. The first category which we call “*fill through copying*” attempts to explicitly borrow contents or textures from surroundings to fill the missing regions. An example is diffusion-based [1, 3] methods which propagate local image appearance surrounding the target holes based on the isophote direction field. Another stream is relying on texture synthesis techniques, which fills the hole by both extending and borrowing textures from surrounding regions [6, 9, 13, 29, 31]. Patch-based algorithms like [9, 10, 11, 29] progressively fill pixels in the hole by searching the image patches from background regions that are the most similar to the pixels along the hole boundaries.

The second group attempts to “*fill through modeling*” and hallucinates missing pixels in a data-driven manner with the use of large external databases. These approaches learn to model the distribution of the training images and assume that regions surrounded by similar contexts likely to possess similar contents [15, 18, 20, 21, 30, 33]. For instance, PixelRNN [20] uses a two-dimensional Recurrent Neural Network (RNN) to model the pixel-level dependencies along two spatial dimensions. More general idea [15, 33] is to train an encoder-decoder convolutional network to model the 2-dimensional spatial contents. Rather than modeling the raw pixels, [18, 30] train a convolutional network to model image-wide edge structure or foreground object contours, thus enabling auto-completion of the edge or contours. These techniques are effective when they find an example image with sufficient visual similarity to the query, but will easily fail if the database does not have similar examples. To overcome the limitation of copying-based or modeling-based methods, the third group of approaches attempts to combine the two [19, 24, 32, 35, 36, 37]. These methods learn to model the image distribution in a data-driven manner, and in the meantime, they develop mechanisms to explicitly borrow patches/features from background regions. [36] introduces a novel contextual attention layer that enables borrowing features from distant spatial locations. [37] further extends the contextual attention mechanism to multiple scales and all the way from feature-level to image-level. [24] employs the patch-swap layer that propagates the high-frequency texture details from the boundaries to hole regions.

Most learning-based approaches belong to the second or

third group. Compared to traditional methods, these techniques have strong ability to learn adaptive and high-level features of disparate semantics and thus are more adept in hallucinating visually plausible contents especially when inpainting structured images like faces [15, 20, 24, 35, 36, 37], objects [18, 21, 30, 33], and natural scenes [15, 24, 35, 36]. Since existing methods employ convolutional layers directly on the original input, the memory usage could become extremely high and intractable when the input size is up to 8K. Another issue is that the quality deteriorates rapidly when hole size increases with image size. Even if the training is feasible, access to large amounts of high-resolution training data would be tedious and expensive.

To resolve these issues, we propose a novel Contextual Residual Aggregation (CRA) mechanism to enable the completion of ultra high-resolution images with limited resources. In specific, we use a neural network to predict a low-resolution inpainted result and up-sample it to yield a large blurry image. Then we produce the high-frequency residuals for in-hole patches by aggregating weighted high-frequency residuals from contextual patches. Finally, we add the aggregated residuals to the large blurry image to obtain a sharp result. Since the network only operates on low-resolution images, the cost of memory and computing time is significantly reduced. Moreover, as the model can be trained with low-resolution images, the need for high-resolution training datasets is alleviated. Furthermore, we introduce other techniques including slim and deep layer configuration, attention score sharing, multi-scale attention transfer, and Light-Weight Gated Convolutions (LWGC) to improve the inpainting quality, computation, and speed. Our method can inpaint images as large as 8K with satisfying quality, which cannot be handled by prior learning-based approaches. Exemplar results are shown in Figure 1. The contributions of the paper are summarized as follows:

- We design a novel and efficient Contextual Residual Aggregation (CRA) mechanism that enables ultra high-resolution inpainting with satisfying quality. The mechanism enables large images (up to 8K) with considerable hole sizes (up to 25%) to be inpainted with limited memory and computing resources, which is intractable for prior methods. Also, the model can be trained on small images and applied on large images, which significantly alleviates the requirements for high-resolution training datasets.
- We develop a light-weight model for irregular hole-filling that can perform real-time inference on images of 2K resolutions on a NVIDIA GTX 1080 Ti GPU, using techniques including slim and deep layer configuration, attention score sharing, and Light Weight Gated Convolution (LWGC).
- We use attention transfer at multiple abstraction levels which enables filling holes by weighted copying features from contexts at multiple scales, improving the

inpainting quality over existing methods by a certain margin even when tested on low-resolution images.

2. Related Works

2.1. Irregular Hole-filling & Modified Convolutions

Vanilla convolutions are intrinsically troublesome for irregular hole-filling because convolutional filters treat all pixels the same as valid ones, causing visual artifacts such as color inconsistency, blurriness, and boundary artifacts. Partial convolution [19] is proposed to handle irregular holes, where the convolution is masked and re-normalized to be conditioned on valid pixels. Gated convolution [35] generalizes the partial convolution idea by providing a learnable dynamic feature selection mechanism for each channel and at each spatial location, achieving better visual performance. Here, we further improve the gated convolution through a lightweight design to improve efficiency.

2.2. Contextual Attention

Contextual attention [36] is proposed to allow long-range spatial dependencies during inpainting, which enables borrowing pixels from distant locations to fill missing regions. The contextual attention layer has two phases: “match” and “attend”. In the “match” phase, the attention scores are computed by obtaining region affinity between patches inside and those outside the holes. In the “attend” phase, holes are filled by copying and aggregating patches from weighted contexts by the attention scores. [37] extends this idea by using a pyramid of contextual attention at multiple layers. In contrast to [37], we only compute the attention scores once and reuse them at multiple abstraction levels, which leads to fewer parameters and less computation.

2.3. Image Residuals

The difference between an image and the blurred version of itself represents the high-frequency image [4, 7]. Early works use the difference obtained by Gaussian blurring for low-level image processing tasks like edge detection, image quality assessment, and feature extraction [8, 23, 26]. We employ this concept to decompose the input image into low-frequency and high-frequency components. The low-frequency component is obtained through averaging neighboring pixels, whereas the high-frequency component (i.e. image residuals) is obtained by subtracting the original image with its low-frequency component.

3. Method

3.1. The Overall Pipeline

Figure 2 illustrates the overall pipeline of the proposed CRA mechanism where the generator is the only trainable component in the framework. Given a high-resolution input image, we first down-sample the image to 512×512 and then up-sample it to obtain a blurry large image of the same size as the raw input (Section 4.1). The height and width of

the image are not necessary to be equal but must be multiples of 512. The generator takes the low-resolution image and fills the holes. Meanwhile, the attention scores are calculated by the Attention Computing Module (ACM) of the generator (Section 3.2.1). Also, the contextual residuals are computed by subtracting the large blurry image from the raw input, and the aggregated residuals in the mask region are then calculated from the contextual residuals and attention scores through an Attention Transfer Module (ATM) (Section 3.2.2). Finally, adding the aggregated residuals to the up-sampled inpainted result generates a large sharp output in the mask region while the area outside mask is simply a copy of the original raw input.

3.2. Contextual Residual Aggregation (CRA)

Filling the missing region by using contextual information [24, 32, 34], and contextual attention mechanism [36] has been proposed previously. Similarly, we deploy the CRA mechanism to borrow information from contextual regions. However, the CRA mechanism borrows from contexts not only features but also residuals. In particular, we adopt the idea of contextual attention [36] in calculating attention scores by obtaining region affinity between patches inside/outside missing regions. Thus contextually relevant features and residuals outside can be transferred into the hole. Our mechanism involves two key modules: Attention Computing Module and Attention Transfer Module.

3.2.1 Attention Computing Module (ACM)

The attention scores are calculated based on region affinity from a high-level feature map (denoted as P in Figure 2). P is divided into patches and ACM calculates the cosine similarity between patches inside and outside missing regions:

$$c_{i,j} = \left\langle \frac{p_i}{\|p_i\|}, \frac{p_j}{\|p_j\|} \right\rangle \quad (1)$$

where p_i is the i^{th} patch extracted from P outside mask, p_j is the j^{th} patch extracted from P inside the mask. Then softmax is applied on the similarity scores to obtain the attention scores for each patch:

$$s_{i,j} = \frac{e^{c_{i,j}}}{\sum_{i=1}^N e^{c_{i,j}}} \quad (2)$$

where N is the number of patches outside the missing hole. In our framework, each patch size is 3×3 and P is 32×32 , thus a total number of 1024 patches can be extracted. In practice, the number of in-hole patches could vary for different hole sizes. We uniformly use a matrix of 1024×1024 to save affinity scores between any possible pair of patches, though only a fraction of them are useful.

3.2.2 Attention Transfer Module (ATM)

After obtaining the attention scores from P , the corresponding holes in the lower-level feature maps (P^l) can be filled with contextual patches weighted by the attention scores:

$$p_j^l = \sum_{i=1}^N s_{i,j} p_i^l \quad (3)$$

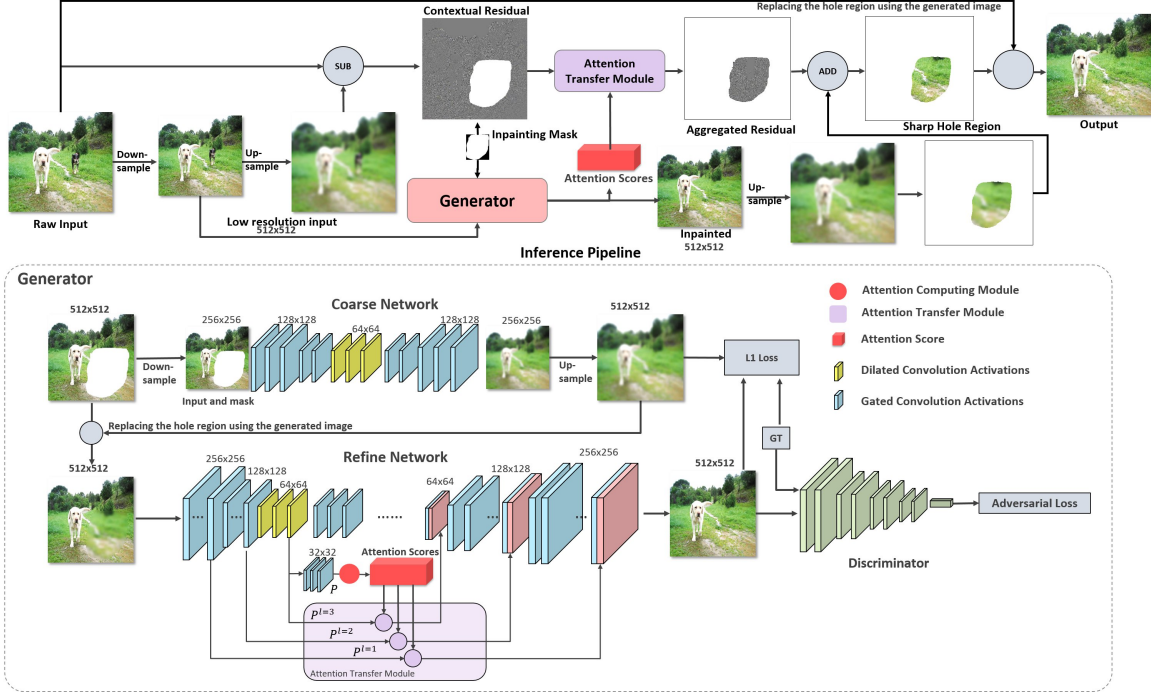


Figure 2: The overall pipeline of the method: (top) CRA mechanism, (bottom) the architecture of the generator.

where $l \in 1, 2, 3$ is the layer number and p_i^l is the i^{th} patch extracted from P^l outside masked regions, and p_j^l is the j^{th} patch to be filled inside masked regions. N indicates the number of contextual patches (background). After calculating all in-hole patches, we can finally obtain a filled feature P^l . As the size of feature maps varies by layer, the size of patches should vary accordingly. Assuming the size of the feature map is 128^2 and the attention scores are computed from 32^2 patches, then the patch sizes should be greater or equal to $(128/32)^2 = 4^2$ so that all pixels can be covered. If the patch sizes are greater than 4×4 , then certain pixels are overlapped, which is fine as the following layers of the network can learn to adapt.

Multi-scale attention transfer and score sharing. In our framework, we apply attention transfer multiple times with the same set of attention scores (Figure 2). The sharing of attention scores leads to fewer parameters and better efficiency in terms of memory and speed.

3.2.3 Residual Aggregation

The target of Residual Aggregation is to calculate residuals for the hole region so that sharp details of the missing contents could be recovered. The residuals for the missing contents can be calculated by aggregating the weighted contextual residuals obtained from previous steps:

$$R_j = \sum_{i=1}^N s_{i,j} R_i \quad (4)$$

where R is the residual image and R_i is the i^{th} patch extracted from contextual residual image outside the mask, and R_j is j^{th} patch to be filled inside the mask. The patch

sizes are properly chosen to exactly cover all pixels without overlapping, to ensure the filled residuals being consistent with surrounding regions. Once the aggregated residual image is obtained, we add it to the up-sampled blurry image of the generator, and obtain a sharp result (Figure 2).

3.3. Architecture of Generator

The network architecture of the generator is shown in Figure 2. We use a two-stage coarse-to-fine network architecture where the coarse network hallucinates rough missing contents, and the refine network predicts finer results. The generator takes an image and a binary mask indicating the hole regions as input and predicts a completed image. The input and output sizes are expected to be 512×512 . In order to enlarge the perceptive fields and reduce computation, inputs are down-sampled to 256×256 before convolution in the coarse network, different from the refine network who operates on 512×512 . The prediction of the coarse network is naively blended with the input image by replacing the hole region of the latter with that of the former as the input to the refine network. Refine network computes contextual attention scores with a high-level feature map and performs attention transfer on multiple lower-level feature maps, thus distant contextual information can be borrowed at multiple abstraction levels. We also adopt dilated convolutions [15] in both coarse and refine networks to further expand the size of the receptive fields. To improve the computational efficiency, our inpainting network is designed in a slim and deep fashion, and the LWGC is applied for all layers of the generator. Other implementation considerations

include: (1) using ‘same’ padding and ELUs [5] as activation for all convolution layers, (2) removing batch normalization layer as they deteriorate color coherency [15].

3.4. Light Weight Gated Convolution

Gated Convolutions (GC) [35] leverages the art of irregular hole inpainting. However, GC almost doubles the number of parameters and processing time in comparison to vanilla convolution. In our network, we proposed three modified versions of GC called Light Weight Gated Convolutions (LWGC), which reduces the number of parameters and processing time while maintaining the effectiveness. The output of the original GC can be expressed as:

$$\begin{aligned} G &= \text{conv}(W_g, I) \\ F &= \text{conv}(W_f, I) \\ O &= \sigma(G) \odot \psi(F) \end{aligned} \quad (5)$$

where σ is Sigmoid function thus the output values are within $[0, 1]$. ψ is an activation function which are set to ELU in our experiments. w_g and w_f are two different set of convolutional filters, which are used to compute the gates and features respectively. GC enables the network to learn a dynamic feature selection mechanism. The three variations of LWGC that we propose are named as: depth-separable LWGC (LWGC^{ds}), pixelwise LWGC (LWGC^{pw}), and single-channel LWGC (LWGC^{sc}). They differ by the computation of the gate branch, G:

$$\begin{aligned} G &= \text{conv}^{\text{depth-separable}}(W_g, I) & (6) \\ G &= \text{conv}^{\text{pixelwise}}(W_g, I) & (7) \\ G &= \text{conv}(W_g, I), \text{ G is single-channel} & (8) \end{aligned}$$

The depth-separable LWGC employs a depth-wise convolution followed by a 1×1 convolution to compute gates. The pixelwise LWGC uses a pixelwise or 1×1 convolution to compute the gates. The single-channel LWGC outputs a single-channel mask that is broadcast to all feature channels during multiplication. The single-channel mask is similar to partial convolution, whereas the mask of partial convolutions is hard-wired and untrainable, and generates a binary mask instead of a soft mask. Given that the height (H_k) and width (W_k) of kernels, and numbers of input channels (C_{in}) and output channels (C_{out}), we compare the number of parameters needed to calculate gates in Table 1. We used the single-channel LWGC for all layers of the coarse network and depth-separable or pixelwise LWGC for all layers of the refine network, which has been proved to be equally effective as regular GC but more efficient (Section 4.2).

3.5. Training of the network

3.5.1 Training Losses

Without the degradation of performance, we also significantly simplify the training objectives as two terms: the adversarial loss and the reconstruction loss. We use the

Table 1: Number of parameters needed to compute gates

Method	Parameters Calculation	$H_k, W_k = 3$ $C_{in}, C_{out} = 32$
GC [35]	$H_k \times W_k \times C_{in} \times C_{out}$	9216
LWGC ^{ds}	$H_k \times W_k \times C_{in} + C_{in} \times C_{out}$	1312
LWGC ^{pw}	$C_{in} \times C_{out}$	1024
LWGC ^{sc}	$H_k \times W_k \times C_{in} \times 1$	288

WGAN-GP loss as our adversarial loss [12], which enforces global consistency in the second-stage refinement network. The discriminator and generator are alternatively trained with the losses defined in Equation 9 and Equation 10:

$$L_d = \mathbb{E}_{\tilde{\mathbf{x}} \in \mathbb{P}_g} [D(\tilde{\mathbf{x}})] - \mathbb{E}_{\mathbf{x} \in \mathbb{P}_r} [D(\mathbf{x})] + \sigma_{\mathbb{E}_{\tilde{\mathbf{x}} \in \mathbb{P}_{\tilde{\mathbf{x}}}} [\|\nabla_{\tilde{\mathbf{x}}} D(\tilde{\mathbf{x}})\|_2 - 1]^2} \quad (9)$$

where $D(\cdot)$ is the discriminator output and $G(\cdot)$ is the generator output. \mathbf{x} , $\tilde{\mathbf{x}}$, $\hat{\mathbf{x}}$, are real images, generated images, and interpolations between them, respectively. $\mathbb{P}_g, \mathbb{P}_r, \mathbb{P}_{\tilde{\mathbf{x}}}$ are the corresponding distributions of them respectively.

$$L_{adv} = -\mathbb{E}_{\tilde{\mathbf{x}} \in \mathbb{P}_g} [D(\tilde{\mathbf{x}})] \quad (10)$$

We also add the $L1$ loss to force the consistency of the prediction with the original image. In contrast to [36], we avoid the computationally expensive spatially-discounted reconstruction loss. For simplicity, we just assign a smaller constant weight for the reconstruction loss of all in-hole pixels. The reconstruction loss is thus written as:

$$L_{in-hole} = |G(\mathbf{x}, \mathbf{m}) - \mathbf{x}| \odot \mathbf{m} \quad (11)$$

$$L_{context} = |G(\mathbf{x}, \mathbf{m}) - \mathbf{x}| \odot (1 - \mathbf{m}) \quad (12)$$

$$L_{rec} = \alpha_1 L_{in-hole} + \alpha_2 L_{context} \quad (13)$$

where α_1 and α_2 are coefficients for the in-hole term and contextual term ($\alpha_1 = 1$, and $\alpha_2 = 1.2$). The coarse network is trained with the reconstruction loss explicitly, while the refinement network is trained with a weighted sum of the reconstruction and GAN losses. The coarse network and refine network are trained simultaneously with merged losses as shown in Equation 14.

$$L_g = L_{rec} + \beta L_{adv} \quad (14)$$

where β is the coefficient for adversarial loss ($\beta = 10^{-4}$).

3.5.2 Random Mask Generation

To diversify the inpainting masks, we use two methods to generate irregular masks on-the-fly during training. The first one is [19], which simulates tears, scratches, spots or manual erasing with brushes. The second approach generates masks by randomly manipulating the real object shape templates, accounting for the object removal scenario. These shape templates are obtained from object segmentation masks and including a wide range of categories such as single, multiple or crowded objects. We also randomly rotate, flip and scale the templates with a random scale ratio. In practice, the aforementioned two methods can be combined or separated, depending on specific needs.




Figure 3: Comparing down-sampling and up-sampling operators: (top) using Bilinear up-sampling along with Averaging down-sampling generates more coherent textures with the surroundings. (bottom) using the Averaging down-sampling along with Nearest Neighbor produces tiling artifacts while Bilinear and Bicubic up-sampling perform equally well.



Figure 4: Comparisons of different Gated Convolutions configurations. For example, the notation of $LWGC^{sc}+LWGC^{ds}$ means: the coarse network uses single-channel LWGC and the refine network uses depth-separable LWGC.

3.5.3 Training Procedure

During training, color values of all images are linearly scaled to $[-1, 1]$ in all experiments, and the mask uses 1 to indicate the hole region and 0 to indicate background. The masked image is constructed as $\mathbf{x} \odot (1 - \mathbf{m})$, where \mathbf{x} is the input image and \mathbf{m} is the binary mask, and \odot represents dot product. Inpainting generator G takes concatenation of the masked image and mask as input, and predicts $\mathbf{y} = G(\mathbf{x}, \mathbf{m})$ of the same size as the input image. The entire training procedure is illustrated in Algorithm 1.

Algorithm 1: Training of our proposed network

```

initialization;
while  $G$  has not converged do
  for  $i = 1, \dots, 5$  do
    Sampling batch images  $\mathbf{x}$  from training data;
    Generating random masks  $\mathbf{m}$  for  $\mathbf{x}$ ;
    Getting inpainted  $\mathbf{y} \leftarrow G(\mathbf{x}, \mathbf{m})$ ;
    Pasting back  $\tilde{\mathbf{x}} \leftarrow \mathbf{y} \odot \mathbf{m} + \mathbf{x} \odot (1 - \mathbf{m})$ ;
    Sampling a random number  $\alpha \in U[0, 1]$ ;
    Getting interpolation  $\hat{\mathbf{x}} \leftarrow (1 - \alpha)\mathbf{x} + \alpha\tilde{\mathbf{x}}$ ;
    Updating the discriminator  $D$  with loss  $L_d$ ;
  end
  Sampling batch images  $\mathbf{x}$  from training data;
  Generating random masks  $\mathbf{m}$  for  $\mathbf{x}$ ;
  Getting Inpainted  $\mathbf{y} \leftarrow G(\mathbf{x}, \mathbf{m})$ ;
  Pasting back  $\tilde{\mathbf{x}} \leftarrow \mathbf{y} \odot \mathbf{m} + \mathbf{x} \odot (1 - \mathbf{m})$ ;
  Updating generator  $G$  with loss  $L_g$ ;
end

```

4. Experimental Results

We evaluate the proposed method on three datasets including Places2 [38], CelebA-HQ [16], and DIV2K [25]. Our model is trained on two NVIDIA 1080 Ti GPUs with images of resolution 512×512 with batch size of 8. For DIV2K and CelebA-HQ, images are down-sampled to 512×512 . For Places2, images are randomly cropped to 512×512 . After training, we test the models on images of various resolutions of 512 to 8K on a GPU. The final model has a total of 2.7M parameters and implemented on TensorFlow v1.13 with CUDNN v7.6 and CUDA v10.0.

4.1. Analysis of CRA Design

As shown in Figure 2, the CRA mechanism involves one down-sampling and two up-sampling operations outside of the generator. Choosing different methods for down-sampling and up-sampling may affect the final results. To explore this, we experimented with four down-sampling methods: Nearest-Neighbor, Bilinear, Bicubic, and Averaging. Averaging evenly splits the input into target patches and average all pixels in each patch to obtain a 512×512 image. We also explored three up-sampling methods including Nearest-Neighbor, Bilinear or Bicubic. Note that the two up-sampling operations must be consistent, so we do not consider inconsistent combinations. Experimental results on an HD dataset indicate that Averaging performs the best for down-sampling and Bilinear or Bicubic performs equally well for up-sampling (Figure 3). For simplicity, we use Averaging down-sampling and Bilinear up-sampling.




Figure 5: Qualitative comparisons using 512×512 (top) and 1024×1024 (bottom) images from Places2 validation dataset.

Table 2: Quantitative evaluation results on Places2 validation set. Note that certain models cause Out-Of-Memory (OOM) error when tested on 2K or 4K images, thus the corresponding cells are left empty.

Image Size	512×512					1024×1024					2048×2048					4096×4096					
	Metrics	L1	MS-SSIM	FID	IS	Time	L1	MS-SSIM	FID	IS	Time	L1	MS-SSIM	FID	IS	Time	L1	MS-SSIM	FID	IS	Time
DeepFillV1[36]	6.733	0.8442	7.541	17.56	62ms	7.270	0.8424	10.21	17.69	663ms	-	-	-	-	-	-	-	-	-	-	-
DeepFillV2[35]	6.050	0.8848	4.939	18.20	78ms	6.942	0.8784	8.347	17.04	696ms	-	-	-	-	-	-	-	-	-	-	-
PEN-Net[37]	9.732	0.8280	14.13	14.51	35ms	10.42	0.8128	19.36	12.51	289ms	-	-	-	-	-	-	-	-	-	-	-
PartialConv[19]	8.197	0.8399	29.32	13.13	35ms	11.19	0.8381	32.20	13.53	110ms	16.19	0.8373	41.23	11.17	920ms	-	-	-	-	-	-
Global-local[15]	8.617	0.8469	21.27	13.48	53ms	9.232	0.8392	26.23	13.05	219ms	9.308	0.8347	27.09	12.61	219ms	-	-	-	-	-	-
Ours	5.439	0.8840	4.898	17.72	25ms	5.439	0.8840	4.899	17.72	31ms	5.492	0.8840	4.893	17.85	37ms	5.503	0.8840	4.895	17.81	87.3ms	-

4.2. Analysis of Light Weight Gated Convolutions

We propose three types of LWGC, which are faster than the original GC. We experimented how they affect inpainting quality and efficiency on the CelebA-HQ dataset to explore the influence of LWGC on the results, by replacing the original GCs with LWGCs for the coarse/refine networks. As shown in Figure 4, the $LWGC^{sc}+LWGC^{sc}$ configuration brings visible artifacts while the other five configurations perform equally well in terms of quality. Considering $LWGC^{sc}+LWGC^{pw}$ requires fewer parameters than the other four, we adopt the $LWGC^{sc}+LWGC^{pw}$ configuration in the generator.

4.3. Comparisons With Learning-based Methods

We compared our method with other state-of-the-art learning-based inpainting methods including Global-local GAN [15], DeepFillV1 [36], DeepFillV2 [35], PEN-Net [37] and Partial Convolution [19]. To make fair comparisons, we attempted to use the same settings for all experiments, though not fully guaranteed. The official pre-trained DeepFillV1 [36] model was trained for 100M iterations with the batch size of 16 and the global-local GAN [15] was trained for 300K iterations with the batch size of 24. Both of them were trained on 256×256 images with rectan-

gular holes of maximum size 128×128 . All the other models were trained for 300K iterations with the batch size of 8 on 512×512 images with irregular holes up to 25% area of the whole image. The original DeepFillV2 [35] model attached a sketch channel to the input to facilitate image manipulation, we simply removed the sketch channel and re-trained the model. For all these methods, no specific post-processing steps are performed other than pasting the filled contents back to the original image.

Qualitative comparisons Figure 5 shows our model performs equally good or slightly better than previous methods on 512×512 images. Partial convolution [19] and global-local GAN [15] performs well when the mask is small and narrow but exert severe artifacts when the hole size becomes bigger. Global-local GAN [15] is problematic in maintaining the color consistency of filled contents with surroundings. DeepFillV1 [36] generates plausible results, but occasionally the artifacts inside the hole region are visible, implying its vulnerability to irregular masks. DeepFillV2 [35] generates incoherent textures when the hole size goes up. Nevertheless, when tested on larger images with bigger hole sizes, our model performs consistently good while the inpainting results of other methods deteriorate dramatically (e.g. severe hole-shaped artifacts in Figure 5).




Figure 6: Comparisons of different super-resolution methods: the red squares area are zoomed-in for more details.




Figure 7: Comparisons of our method with Inpaint (software), Photoshop content-aware fill and an open-source PatchMatch implementation [28]. The masks for Photoshop and Inpaint are manually drawn, thus not guaranteed to be the same.

Quantitative comparisons Table 2 reports our quantitative evaluation results in terms of mean $L1$ error, MS-SSIM [27], Inception Score (IS) [22], and Frechet Inception Distance (FID) [14]. It also shows the average inference time per image on a NVIDIA GTX 1080 Ti GPU. These metrics are calculated over all 36,500 images of the Places2 validation set. Each image is assigned a randomly generated irregular mask. To examine the performance on various image sizes, we linearly scale images and masks to various dimensions. Table 2 shows that our proposed model achieves the lowest $L1$ loss and FID on 512×512 images. When the input images are greater than or equal to 1024×1024 , our proposed model achieves the best results on all metrics. In addition, the proposed approach significantly outperforms other learning-based methods in terms of speed. In specific, it is 28.6% faster for 512×512 images, 3.5 times faster for 1024×1024 images, and 5.9 times faster for 2048×2048 images than the second-fastest method. Furthermore, the proposed model can inpaint 4096×4096 images in 87.3 milliseconds which is intractable with other learning-based methods due to limits of GPU memory.

4.4. Comparisons of CRA with Super-resolution

Figure 6, compares the high-resolution results of our CRA with those obtained via various super-resolution techniques. After obtaining a 512×512 inpainted result, we up-sample the output to the original size using different up-sampling methods including SRGAN [17], Nearest Neighbor, and Bicubic, then, we paste the filled contents to the original image. SRGAN [17] is a learning-based method that can perform $4 \times$ super-resolution and the official pre-trained model was trained on DIV2K. We can observe from that the hole region generated by CRA is generally sharper

and visually more consistent with surrounding areas.

4.5. Comparisons With Traditional Methods

Moreover, we compare our method with two commercial products based on PatchMatch [2] (Photoshop, Inpaint) and an open-source PatchMatch implementation for image inpainting [28] (Figure 7). We discover that PatchMatch-based methods are able to generate clear textures but with distorted structures incoherent with surrounding regions.

5. Conclusion

We presented a novel Contextual Residual Aggregated technique that enables more efficient and high-quality inpainting of ultra high-resolution images. Unlike other data-driven methods, the increase of resolutions and hole size does not deteriorate the inpainting quality and does not considerably increase the processing time in our framework. When tested on high-resolution images between 1K and 2K, our model is extremely efficient where it is $3x \sim 6x$ faster than the state-of-the-art on images of the same size. Also, it achieves better quality by reducing FID by 82% compared to the state-of-the-art. We also compared our method with commercial products that showed significant superiority in certain scenarios. So far, our method is the only learning-based technique that can enable end-to-end inpainting on the ultra-high-resolution image ($4K \sim 8K$). In the future, we will explore similar mechanisms for other tasks such as image expansion, video inpainting and image blending.

Acknowledgment

We want to acknowledge Peng Deng, Shao Hua Chen, Xinjiang Sun, Chunhua Tian and other colleagues in Huawei Technologies for their support in the project.

References

- [1] Coloma Ballester, Marcelo Bertalmio, 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):1200–1211, 2001.
- [2] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. In *ACM Transactions on Graphics (ToG)*, volume 28, page 24. ACM, 2009.
- [3] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 417–424. ACM Press/Addison-Wesley Publishing Co., 2000.
- [4] Peter Burt and Edward Adelson. The laplacian pyramid as a compact image code. *IEEE Transactions on communications*, 31(4):532–540, 1983.
- [5] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- [6] Antonio Criminisi, Patrick Pérez, and Kentaro Toyama. Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on image processing*, 13(9):1200–1212, 2004.
- [7] Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pages 1486–1494, 2015.
- [8] Manjusha Deshmukh and Udhav Bhosale. Image fusion and image quality assessment of fused images. *International Journal of Image Processing (IJIP)*, 4(5):484, 2010.
- [9] Iddo Drori, Daniel Cohen-Or, and Hezy Yeshurun. Fragment-based image completion. In *ACM Transactions on graphics (TOG)*, volume 22.3, pages 303–312. ACM, 2003.
- [10] Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 341–346. ACM, 2001.
- [11] 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.
- [12] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017.
- [13] Kaiming He and Jian Sun. Statistics of patch offsets for image completion. In *European Conference on Computer Vision*, pages 16–29. Springer, 2012.
- [14] 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 *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017.
- [15] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):107, 2017.
- [16] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [17] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [18] Liang Liao, Ruimin Hu, Jing Xiao, and Zhongyuan Wang. Edge-aware context encoder for image inpainting. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3156–3160. IEEE, 2018.
- [19] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 85–100, 2018.
- [20] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016.
- [21] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [22] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- [23] Mohsen Sharifi, Mahmood Fathy, and Maryam Tayefeh Mahmoudi. A classified and comparative study of edge detection algorithms. In *Proceedings. International conference on information technology: Coding and computing*, pages 117–120. IEEE, 2002.
- [24] Yuhang Song, Chao Yang, Zhe Lin, Xiaofeng Liu, Qin Huang, Hao Li, and C-C Jay Kuo. Contextual-based image inpainting: Infer, match, and translate. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018.
- [25] Radu Timofte, Shuhang Gu, Jiqing Wu, Luc Van Gool, Lei Zhang, Ming-Hsuan Yang, Muhammad Haris, et al. Ntire 2018 challenge on single image super-resolution: Methods and results. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2018.
- [26] Alexander Toet. Image fusion by a ratio of low-pass pyramid. *Pattern Recognition Letters*, 9(4):245–253, 1989.
- [27] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multi-scale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003.
- [28] Zubeyr Furkan Eryilmaz Wen-Fu Lee, Yuan-Ting Hsieh. image completion using patchmatch algorithm. https://github.com/YuanTingHsieh/Image_Completion. Accessed: 2019-10-26.

- [29] Marta Wilczkowiak, Gabriel J Brostow, Ben Tordoff, and Roberto Cipolla. Hole filling through photomontage. In *BMVC 2005-Proceedings of the British Machine Vision Conference 2005*, 2005.
- [30] Wei Xiong, Jiahui Yu, Zhe Lin, Jimei Yang, Xin Lu, Connelly Barnes, and Jiebo Luo. Foreground-aware image inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5840–5848, 2019.
- [31] Zongben Xu and Jian Sun. Image inpainting by patch propagation using patch sparsity. *IEEE transactions on image processing*, 19(5):1153–1165, 2010.
- [32] Zhaoyi Yan, Xiaoming Li, Mu Li, Wangmeng Zuo, and Shiguang Shan. Shift-net: Image inpainting via deep feature rearrangement. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 1–17, 2018.
- [33] Chao Yang, Xin Lu, Zhe Lin, Eli Shechtman, Oliver Wang, and Hao Li. High-resolution image inpainting using multi-scale neural patch synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6721–6729, 2017.
- [34] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
- [35] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. *arXiv preprint arXiv:1806.03589*, 2018.
- [36] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5505–5514, 2018.
- [37] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Baining Guo. Learning pyramid-context encoder network for high-quality image inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1486–1494, 2019.
- [38] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017.