

Learning Oracle Attention for High-fidelity Face Completion

Tong Zhou¹ Changxing Ding¹ Shaowen Lin¹ Xinchao Wang² Dacheng Tao³

¹ South China University of Technology ² Stevens Institute of Technology

³ UBTECH Sydney AI Centre, School of Computer Science, Faculty of Engineering,
The University of Sydney, Darlington, NSW 2008, Australia

201821011282@mail.scut.edu.cn chxding@scut.edu.cn

eeswlin@mail.scut.edu.cn xinchao.w@gmail.com dacheng.tao@sydney.edu.au

Abstract

High-fidelity face completion is a challenging task due to the rich and subtle facial textures involved. What makes it more complicated is the correlations between different facial components, for example, the symmetry in texture and structure between both eyes. While recent works adopted the attention mechanism to learn the contextual relations among elements of the face, they have largely overlooked the disastrous impacts of inaccurate attention scores; in addition, they fail to pay sufficient attention to key facial components, the completion results of which largely determine the authenticity of a face image. Accordingly, in this paper, we design a comprehensive framework for face completion based on the U-Net structure. Specifically, we propose a dual spatial attention module to efficiently learn the correlations between facial textures at multiple scales; moreover, we provide an oracle supervision signal to the attention module to ensure that the obtained attention scores are reasonable. Furthermore, we take the location of the facial components as prior knowledge and impose a multi-discriminator on these regions, with which the fidelity of facial components is significantly promoted. Extensive experiments on two high-resolution face datasets including CelebA-HQ and Flickr-Faces-HQ demonstrate that the proposed approach outperforms state-of-the-art methods by large margins.

1. Introduction

Image inpainting refers to filling in the missing pixels in an image with the expectation that the recovered image will be visually realistic. This process not only requires that the filled textures themselves be meaningful, but also seeks semantic consistency between the filled area and the context. Image inpainting is widely applied to photo restoration, image editing, and object removal, among many other tasks.




Figure 1. Face completion results for images (256×256) with center mask (128×128). In each row from left to right: (a) the result taken from the paper of PEN-Net [34]. By zooming in, we can observe the color discrepancy between two eyes, along with distortions in the nose and mouth regions. (b) The results of our method. It can be seen that our method indeed produces face images with high fidelity.

Face completion, as a branch of image inpainting, focuses on filling the missing regions of a human face, and turns out to be a challenging task. The reasons behind lie in two aspects. First, the human face contains rich and subtle textures that also differ dramatically across persons, mean-

ing that it is difficult to perfectly restore these diverse facial textures. Second, there are close correlations between facial components, making the fidelity of the image more vulnerable to the semantic consistency between the facial components. Take one very recent work [34] for example, its generated images are satisfactory in facial structure, but still suffer from small artifacts in facial components and semantic inconsistency. As illustrated in Figure 1, the two eyes in the same face are different in color; moreover, small distortions in the nose and mouth areas can also be observed. These flaws have a substantial impact on the authenticity of the overall visual effect.

Recently, convolutional neural network (CNN) based methods have become the mainstream approach to image inpainting [30, 29, 12, 26, 36]. In order to generate visually realistic images, existing approaches typically adopt the non-local schemes that make use of contextual relations to fill in the missing pixels [32, 21, 39, 24]. However, due to the lack of direct supervision on the obtained attention scores, the learned relations are insufficiently reliable, meaning that these methods may generate distorted textures.

Moreover, different types of structural information have been extracted by networks in order to act as prior knowledge for assisting image inpainting; for example, segmentation maps [25], object contours [28], edge [17], facial landmarks [23], and face parsing [13]. While these methods focus on the correctness of the structural information, they also ignore the quality of the textures on key areas in the image (e.g. facial components in face images).

In this paper, we propose a comprehensive framework to handle the above issues. Inspired by recent progress in attention models [35, 3], we propose a Dual Spatial Attention (DSA) model that comprises foreground self-attention and foreground-background cross-attention modules. To capture contextual information across different scales, we apply it to multiple layers in the network. Compared with attention models introduced in recent works [32, 15], DSA has two key advantages: first, it is more efficient and can capture more comprehensive contextual information; second, we impose an oracle supervision signal to ensure that the attention scores produced by DSA are reasonable. With the help of DSA, our approach obtains semantically consistent face completion results, as illustrated in Figure 1.

Furthermore, we also extract facial landmarks from the ground truth image to act as prior knowledge. Rather than imposing constraints to ensure that the facial landmarks between the recovered and ground-truth images are consistent, we use facial landmarks to locate four key facial components: i.e., both eyes, the nose, and the mouth. Subsequently, we train four discriminators for each recovered facial component. By benefiting from adversarial learning on the specified locations, our generator pays more attention to the textures of each key facial component. As a result,

our proposed approach can generate more visually realistic textures, as shown in Figure 1. Since all discriminators are removed during testing, our approach does not yield any drop in efficiency.

We conduct a number of experiments on high-resolution face datasets, including CelebA-HQ [8] and Flickr-Faces-HQ [9]. Quantitative and qualitative comparisons demonstrate that our proposed approach outperforms state-of-the-art methods by large margins.

2. Related Work

2.1. Image Inpainting

Previous image inpainting methods can be divided into two categories: hand-crafted methods and learning-based methods.

Methods in the first category attempt to copy similar patches from the unmasked area to fill in the missing area [37, 20, 11, 5]. Criminisi *et al.* [2] established the priority of each missing patch with reference to the surrounding structure. When the missing patch is surrounded by more valid pixels or closer to boundary region, this patch is assigned a higher priority. These methods can restore a continuous structure as they give priority to well-defined structural areas; however, iterating over every patch to search for the most similar one results in elevated time and memory costs. Accordingly, Barnes *et al.* [1] proposed a faster method referred to as PatchMatch. This approach employs a randomized algorithm to quickly find the approximate nearest neighbor matches. However, the hand-crafted methods are unable to handle complex-structured or largely occluded images, as only low-level features are considered.

The second category of methods [31, 19, 27] typically involve training a deep CNN in an encoder-decoder structure to predict each pixel of the missing area. Pathak *et al.* [18] proposed Context Encoders that applied adversarial learning to the entire image. In order to generate more realistic details, moreover, Iizuka *et al.* [6] appended an extra local discriminator to improve the generated effect of the masked area. However, this approach relies on post-processing to alleviate artifacts. To address this problem, Liu *et al.* [14] used only valid pixels per convolution to alleviate artifacts by updating a binary mask to indicate the generated pixels.

Face Completion Due to the complexity and diversity of facial structures, face completion is one of the more challenging image inpainting tasks. Generally speaking, researchers in this area therefore use a wealth of face prior knowledge to aid in restoration. For example, Li *et al.* [13] used face parsing to propose a semantic parsing loss. Song *et al.* [23] trained an extra network to restore the facial landmarks and face parsing, then input them along with the masked image to train face completion network. However, the results of these methods are greatly affected

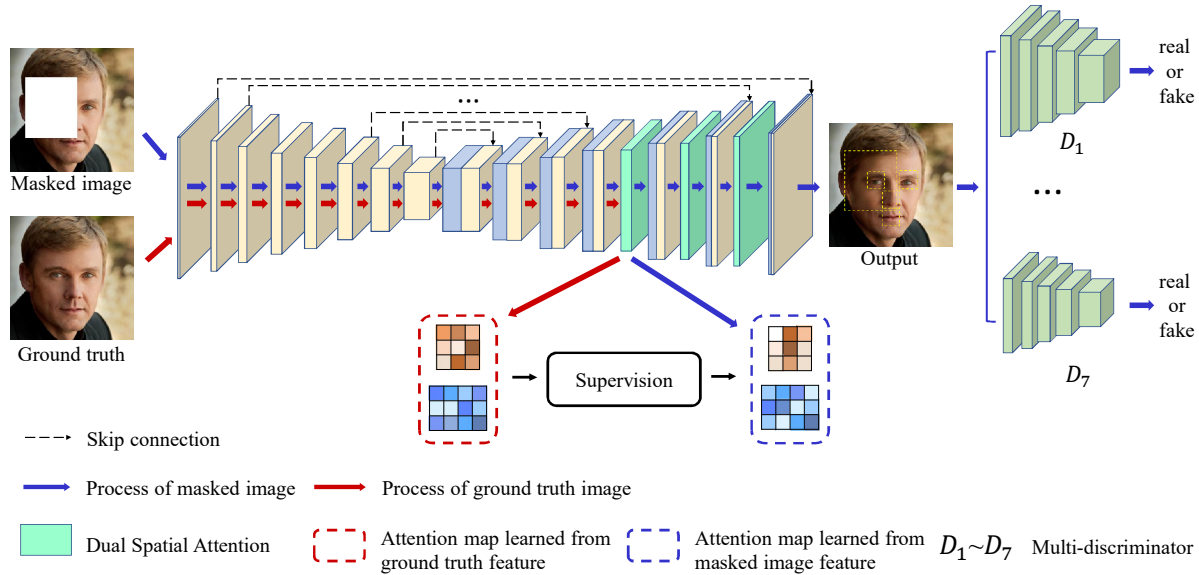


Figure 2. The overall architecture of our model. We use the U-Net structure [14] as backbone. The proposed attention module is embedded to layer 12, 13 and 14, where the resolution is 16×16 , 32×32 and 64×64 respectively. We feed the ground truth images to the network, such that we can impose an oracle supervision signal on attention scores produced by DSA. We deploy discriminators on both the masked area and each of the four facial components, respectively. Best viewed in color.

by the performance of the prior knowledge extracting network, which may consume a large amount of computational resources; moreover, these methods cannot directly guide the network to focus on the texture of key facial components.

2.2. Attention Mechanism

In order to maintain contextual consistency, Yu *et al.* [32] proposed a coarse-to-fine network containing a contextual attention module that learns the correlations between the missing and unmasked patches. Subsequently, some methods directly use the coarse-to-fine network with contextual attention; for example, GConv [33] turns the binary mask proposed by PConv [14] into learnable soft value as a gating mechanism, while Xiong *et al.* [28] used object contours as prior knowledge to assist in restoration. On the other hand, other methods opt to use contextual attention in different ways. For example, Sagong *et al.* [21] designed a parallel decoding network to replace the coarse-to-fine structure, thereby reducing the number of parameters, while Zeng *et al.* [34] proposed an attention transfer network that uses the attention learned from high-level features to reconstruct low-level features. Moreover, the CSA layer [15] learns correlations between the patches inside the masked area.

In summary, previous attention-based methods learn long-range correlations in order to search the similar feature-level patches as references for filling. However, the learned attention is not reliable enough, as the parameters of the attention module lack direct supervision.

3. Proposed Algorithm

The overall architecture of the proposed approach is illustrated in Figure 2. We adopt the same U-Net structure as used in [14] to construct the basic generator. In what follows, we provide the details of our proposed methods. Specifically, we introduce the proposed DSA module with an oracle supervision signal, and then describe the deployment of the multi-discriminator. Finally, we describe the loss functions used to guide the training process.

3.1. Dual Spatial Attention with Supervision

We treat the masked area as the foreground and the unmasked area as background. When learning relations between the different parts of the face, we consider two key scenarios. First, when restoring the foreground features, we obtain reference information from the background. For example, when the left eye is masked and the right one is not, we obtain features from the right eye to help restore the left eye. Second, when the masked area is large, we consider the relations between different parts in the foreground. For instance, when both eyes are masked, we ensure that the restored eyes have similar features.

Inspired by the principle of self-attention [35, 3], we propose the DSA module, which comprises foreground self-attention and foreground background cross-attention modules, so as to tackle above two scenarios. As shown in Figure 3, we first use the mask to segment the input feature \mathbf{X} into the foreground feature \mathbf{X}_p and background feature

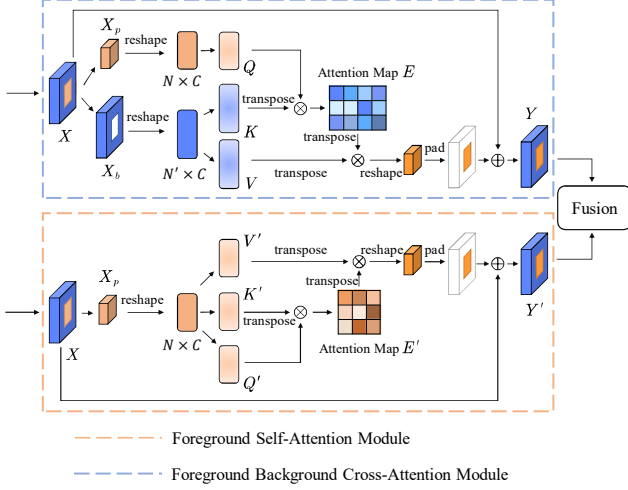


Figure 3. Overview of the Dual Spatial Attention (DSA) module, which includes two parallel branches. They focus on learning foreground-background cross-attention and foreground self-attention, respectively. Best viewed in color.

\mathbf{X}_b . Then we reshape \mathbf{X}_p into $N \times C$ and \mathbf{X}_b into $N' \times C$, where N and N' denote the number of foreground pixels and background pixels in \mathbf{X} respectively. C denotes the number of channels. For the foreground background cross-attention module, we put \mathbf{X}_p and \mathbf{X}_b into 1-dimensional convolutional layers respectively and generate two new feature maps \mathbf{Q} and \mathbf{K} . Subsequently, we conduct a matrix multiplication between \mathbf{Q} and the transpose of \mathbf{K} . Then we apply a softmax layer to obtain the attention map with size of $N \times N'$. We write,

$$E_{ij} = \frac{\exp(\mathbf{Q}_i \cdot \mathbf{K}_j)}{\sum_{n=1}^{N'} \exp(\mathbf{Q}_i \cdot \mathbf{K}_n)}, \quad (1)$$

where E_{ij} denotes the degree of correlation between the i^{th} feature vector of \mathbf{Q} (learned from \mathbf{X}_p) and the j^{th} feature vector of \mathbf{K} (learned from \mathbf{X}_b).

Meanwhile, we also feed feature \mathbf{X}_b into the other 1-dimensional convolutional layer, so as to generate a new feature \mathbf{V} . Next, we perform a matrix multiplication between the transpose of \mathbf{V} and the transpose of the attention matrix, then reshape it back to the original size of \mathbf{X}_p ($C \times H \times W$). In this way, the original foreground features are eventually rebuilt with consideration given to the correlation with the background features. Finally, we extend the rebuilt feature to the original feature map size by means of zero-padding, and then merge it into the original feature map X via element-wise sum operation, which can be formulated as follows:

$$\mathbf{Y} = \alpha \text{Pad}(\mathbf{V}^T \otimes \mathbf{E}^T) + \mathbf{X}, \quad (2)$$

where Pad denotes zero-padding operation, α is a training

parameter that is initialized to zero, and \otimes denotes matrix multiplication.

For the foreground self-attention module, we follow the same steps as for the foreground background cross-attention module, except that we use \mathbf{X}_p only. Specifically, we use three convolutional layers to generate feature maps \mathbf{V}' , \mathbf{K}' , and \mathbf{Q}' . The foreground attention matrix can be formulated as

$$E'_{ij} = \frac{\exp(\mathbf{Q}'_i \cdot \mathbf{K}'_j)}{\sum_{n=1}^N \exp(\mathbf{Q}'_i \cdot \mathbf{K}'_n)}, \quad (3)$$

where E'_{ij} denotes the degree of correlation between the i^{th} feature vector of \mathbf{Q}' and j^{th} feature vector of \mathbf{K}' , where \mathbf{Q}' and \mathbf{K}' are both learned from \mathbf{X}_p . The output can thus be formulated as follows,

$$\mathbf{Y}' = \beta \text{Pad}(\mathbf{V}'^T \otimes \mathbf{E}'^T) + \mathbf{X}, \quad (4)$$

where β is a training parameter initialized to zero as α . Finally, we fuse \mathbf{Y} and \mathbf{Y}' by means of element-wise addition, and then adopt a convolutional layer to adjust the feature and obtain the final refined feature map.

Supervision Signal The attention module helps the network select reference features to improve the filling quality. Nevertheless, if the learned attention is insufficiently accurate, the network may refer to unsuitable features, resulting in poorer filling quality. As a result, ensuring the accuracy of the learned attention scores is key to improving the filling quality. However, the parameters of the attention module are optimized together with the total face completion network using supervision signals for face completion only, meaning that the parameters of the attention module lack a direct restriction.

In order to impose a direct supervision signal towards attention, we extract the attention learned from the ground truth images as the objective. More specifically, in addition to taking the masked image as input, we also feed the ground truth image into the network when training. Through the use of the same network layers, including the DSA, we can obtain the oracle attention maps learned from the ground truth denoted as \mathbf{E}^{gt} and \mathbf{E}'^{gt} . Moreover, all attention scores are received through the softmax layer, which means that the attention map comprises a number of probability distributions. Therefore, we use KL-divergence distance to set up the objective function for attention. The KL-divergence loss is formulated as follows:

$$\begin{aligned} \mathcal{L}_{\text{KL}} = & \frac{1}{N \cdot N'} \sum_{i=1}^N \sum_{j=1}^{N'} E_{ij}^{gt} \cdot (\ln E_{ij}^{gt} - \ln E_{ij}) \\ & + \frac{1}{N \cdot N} \sum_{i=1}^N \sum_{j=1}^N E'_{ij}{}^{gt} \cdot (\ln E'_{ij}{}^{gt} - \ln E'_{ij}), \end{aligned} \quad (5)$$

where we add the average KL-divergence distance between \mathbf{E}^{gt} and \mathbf{E} , as well as between \mathbf{E}'^{gt} and \mathbf{E}' .

In addition, attentions learned from different scale layers complement each other. On a high-resolution feature map, the attention reflects the relations between small-scale features, such as hair. Conversely, the attention on a low-resolution feature map reflects the relations between large-scale features, such as large-range aspects of facial structure. After accounting for the trade-off between computational efficiency and attention learning, we embed the DSA module into three layers of the decoder (*i.e.* layers 12, 13, and 14) to assist the filling process. It is important to note that the large-scale attention represents more large-range structural information and can affect later features; thus, we only impose the KL-divergence loss on layer 12 in order to strengthen the authenticity of the face structure, which can also benefit the subsequent attention learning.

Discussion Compared with the contextual attention layer [32], our DSA uses matrix multiplication rather than dividing the background patches as kernels for convolution, which significantly promotes efficiency. In addition, the CSA layer [15] also learns the relations between foreground patches. However, it computes the similarity between two foreground patches in an iterative way; this slows down the filling speed, especially when the masked area is large. Moreover, neither contextual attention or CSA ensures the accuracy of learned attention scores, which leads to unsatisfying results.

3.2. Multi-discriminator Design

Adversarial learning is helpful in generating photo-realistic images by training the generator and discriminator until a Nash equilibrium is achieved. In addition to the global discriminator [18], which is used on the whole image, the local discriminator is designed to focus on the generated details of the masked area [6, 13]. For face completion tasks, the quality of the facial components largely determines the authenticity of the entire face image. However, relying solely on the global and local discriminators is insufficient if our goal is to guide the network to focus on small regions. To this end, we propose the multi-discriminator for enhancing face details especially on facial components, as shown in Figure 4.

We first use the facial landmarks extracted by the method [10] to mark the locations of the left eye, right eye, nose, and mouth. We then generate four masks of fixed size, and use these masks to crop four facial components. During training, we input each facial component of the generated image and ground truth into the corresponding discriminator to judge whether it is real or fake. Moreover, inspired by the collocation of the global and local discriminators, we further divide the masked area into four equal parts without overlap. These four parts share a discriminator, named the




Figure 4. Overview of the multi-discriminator. G denotes the generator. The four facial components are located by facial landmarks of the ground truth image. Best viewed in color.

local subdivision discriminator, which focuses on more detailed characteristics and can also be regarded as a weighted optimization of the local discriminator loss function.

By using the multi-discriminator, the generator can learn more specific features for each facial component and further enhance the details within the masked area. Unlike previous ways that utilize the prior knowledge, we only use facial landmarks to mark the positions of facial components, after which we guide the network to improve the details in specific areas. Furthermore, since the discriminator only works during the training process, adding multiple discriminators does not affect the efficiency of the implementation.

3.3. Loss Functions

In order to effectively guide the training process, it is critical to design loss functions to be capable of measuring the distance between the generated image and the corresponding ground truth. We thus adopt multiple loss functions from different aspects, described as follows.

Given an original ground truth image \mathbf{I}_{gt} and a randomly generated binary mask \mathbf{M}_1 (zero for holes), we produce the training image \mathbf{I}_{in} by means of element-wise multiplication and denote the generated area of output \mathbf{I}_{out} as \mathbf{I}_{gen} . Moreover, in order to crop the facial components, we also produce a left-eye mask, right-eye mask, nose mask, and mouth mask, which are denoted by $\mathbf{M}_2, \dots, \mathbf{M}_5$ respectively.

Firstly, we use L1-distance between the output and the ground truth as the reconstruction loss \mathcal{L}_r to constrain the value of pixels. Furthermore, we propose to increase the penalty on the facial components and masked area. \mathcal{L}_r can be represented as

$$\mathcal{L}_r = \left\| \left(\mathbf{1} + \sum_{i=1}^5 (\mathbf{1} - \mathbf{M}_i) \right) \odot (\mathbf{I}_{out} - \mathbf{I}_{gt}) \right\|_1, \quad (6)$$

where \odot denotes element-wise multiplication, and $\mathbf{1}$ indicates a matrix with all values set to 1. Here, the facial components and missing area have more weight than other parts,

Method	L1	PSNR	SSIM	LPIPS [38]
PM [1]	5.82%	17.60	0.7786	0.2221
CA [32]	1.82%	24.58	0.8980	0.0977
PIC [39]	1.81%	25.31	0.9023	0.0897
GConv [33]	1.89%	26.29	0.8996	0.0809
Ours	1.46%	26.36	0.9107	0.0706

Table 1. Quantitative results on the same test set using masks with random position. Higher SSIM and PSNR values are better; lower L1 error and LPIPS are better.

while the missing portion of facial components is assigned the largest weight. Secondly, we introduce the perceptual loss L_p using a VGG-16 [22], which is pre-trained by [38] to impose a constraint:

$$\mathcal{L}_p = \|\Psi(\mathbf{I}_{out}) - \Psi(\mathbf{I}_{gt})\|_1 + \|\Psi(\mathbf{I}_{gen}) - \Psi(\mathbf{I}_{gt})\|_1, \quad (7)$$

where Ψ is the output of the pre-trained VGG-16. The perceptual loss computes the L1-distance in feature space between both I_{out} and I_{gen} and the ground truth.

Thirdly, we adopt PatchGAN [7] as our discriminator structure, which maps the input image to a matrix where each element represents the authenticity of a portion of the input image. In this way, the network pays more attention to the local image details. In addition, we adopt an improved version of WGAN with a gradient penalty term [4]. The final adversarial loss function for each discriminator is as follows:

$$\begin{aligned} \mathcal{L}_{D_i} = & \mathbb{E}_{\mathbf{I}_{out} \sim \mathbb{P}_g} [D_i(C_i(\mathbf{I}_{out}))] \\ & - \mathbb{E}_{\mathbf{I}_{gt} \sim \mathbb{P}_{data}} [D_i(C_i(\mathbf{I}_{gt}))] \\ & + \gamma \mathbb{E}_{\hat{\mathbf{I}} \sim \mathbb{P}_1} (\|\nabla_{\hat{\mathbf{I}}} D_i(C_i(\hat{\mathbf{I}}))\|_2 - 1)^2, \end{aligned} \quad (8)$$

where $D_i (i = 1, \dots, 7)$ indicates one of the seven discriminators as shown in Figure 4. Here, C_i denotes the crop operation used to obtain the corresponding areas from the image, $\hat{\mathbf{I}}$ is interpolated from pairs of points sampled from the real data distribution \mathbb{P}_{data} and the generated distribution \mathbb{P}_g ; ∇ means the operation to compute gradient, and γ is set to 10. Thus, the adversarial loss for the generator is as follows,

$$\mathcal{L}_G = - \sum_{i=1}^7 \mathbb{E}_{\mathbf{I}_{out} \sim \mathbb{P}_g} [D_i(C_i(\mathbf{I}_{out}))]. \quad (9)$$

The KL-divergence loss \mathcal{L}_{KL} is depicted as Equation 5, while the details are described in Section 3.1. We define the overall loss function as follows,

$$\mathcal{L} = \lambda_r \mathcal{L}_r + \lambda_{KL} \mathcal{L}_{KL} + \lambda_p \mathcal{L}_p + \lambda_G \mathcal{L}_G, \quad (10)$$

where we empirically set the four trade-off parameters λ_r , λ_{KL} , λ_p , and λ_G to 10, 2, 1 and 1 respectively.

Method	L1	PSNR	SSIM	LPIPS [38]
PEN-Net [34]	2.87%	24.53	0.8369	0.1701
Ours	2.28%	26.11	0.8718	0.1355

Table 2. Quantitative comparison on the same test set with a center mask. Higher SSIM and PSNR values are better; lower L1 error and LPIPS are better.

4. Experiment

4.1. Experimental settings

Dataset We conduct a number of experiments on two high-quality human face datasets including CelebA-HQ [8] and Flickr-Faces-HQ [9]; these contain 30,000 and 70,000 high-quality face images respectively with size of 1024×1024 . We randomly selected 2,000 images from CelebA-HQ and 10,000 from Flickr-Faces-HQ to comprise the test set.

Implementation Details We resize all input images to 256×256 . The random position rectangle masks account for 13.5% \sim 25% of the original image; the largest size is 128×128 and the smallest size is 94×94 . During training, we use RMSprop as the optimizer with a learning rate of 0.0001. On a single NVIDIA TITAN Xp (12GB), we train our model on CelebA-HQ for four days and Flickr-Faces-HQ for eight days with a batch size of 16.

4.2. Comparisons with State-of-the-Art Methods

We conduct qualitative and quantitative comparisons with multiple methods including PatchMatch (PM) [1], Contextual Attention (CA) [32], PIC [39], GConv [33] and PEN-Net [34]. We use the officially released models of CA, PIC, and GConv trained on CelebA-HQ to facilitate fair comparisons. We compare with PEN-Net on the same test set on CelebA-HQ as [34]. As PEN-Net only handles a center mask on CelebA-HQ whose size is 128×128 , we compare with it separately using a center mask. For PIC, we follow the official instructions to select the best of their multiple results. In addition, we train models of CA, PIC, and GConv for Flickr-Faces-HQ respectively. The results on Flickr-Faces-HQ are presented in the supplementary material. In order to focus on assessing the generative capability of the different models, we copy the valid pixels to the output image for all models in comparison.

Quantitative Comparisons As outlined in Table 1 and Table 2, we conduct quantitative comparisons on CelebA-HQ using different masks. We select the commonly-used L1 loss, peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) as evaluation metrics at pixel space. As mentioned in [32, 16], however, these classical metrics are not optimal for image inpainting tasks; thus, we further use Learned Perceptual Image Patch Similarity (LPIPS) [38] as perceptual metric. According to these metrics, our method




Figure 5. Comparisons in visual effect by different methods with random rectangle masks. Four state-of-the-art methods are compared: PM [1], CA [32], PIC [39] and GConv [33]. Best viewed with zoom-in and pay attention to the details on facial components. More qualitative results are presented in the supplementary material.

Method	L1	PSNR	SSIM	LPIPS	Time Cost
CA [32]	1.64%	25.56	0.9057	0.0899	12.9ms
DSA	1.46%	26.36	0.9107	0.0706	6.4ms

Table 3. Quantitative comparison results on CelebA-HQ. The first row is the result of our model that replaces DSA with the CA module [32]. Higher SSIM and PSNR values are better; lower L1 error and LPIPS are better. We compare the time cost per image by CA and DSA on layer 14 of the U-Net.

not only outperforms the previous state-of-the-art image inpainting methods, but also achieves a notable improvement. This is because our method focuses on learning the exact relations of the facial structure and improving the completion effect of facial components; by contrast, other methods tend to ignore the complexity and diversity of the facial structure.

Qualitative Comparisons As shown in Figure 1 and Figure 5, we respectively compare the visual effect using a center mask and a random mask on the CelebA-HQ dataset. Firstly, the hand-crafted method PM [1] fails to recover the basic structure of the face because it is difficult to find similar patches on background. Secondly, CA [32], PIC [39], GConv [33] and PEN-Net [34] employ attention modules to learn contextual information for inpainting. But they still generate semantically inconsistent structures or textures as their learned attention scores may not be reasonable. In addition, these methods may generate artifacts or blurry effect on facial components such as the noses in the first row of Figure 5. This is because they do not pay enough at-

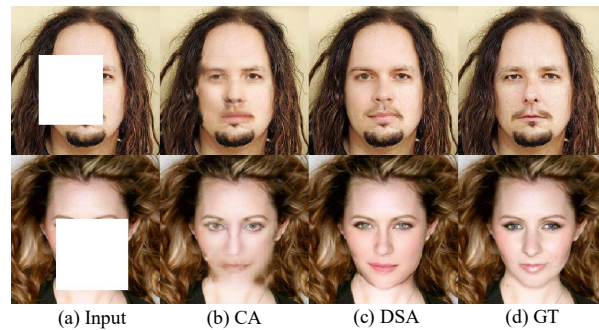


Figure 6. Qualitative comparison results on CelebA-HQ. In each row from left to right: (a) input image, (b) the result of using CA module [32] to replace DSA layer, (c) the result of using DSA module and (d) the ground truth image. Best viewed with zoom-in.

tention to facial components. Finally, our approach produces sharper and more abundant details, especially for facial components, which can be explained by the effect of multi-discriminator. More qualitative results are presented in the supplementary material.

4.3. Ablation Study

Effect of DSA Module We replace the DSA module in our model with the contextual attention (CA) module [32] for comparison. As shown in Table 3 and Figure 6, the model with DSA is more effective and produces images with fine-grained textures. This is because DSA explores pixel-to-pixel correlations rather than patch-to-patch correlation-




Figure 7. Qualitative comparison results on CelebA-HQ. In each row from left to right: (a) input image, (b) the result of using global and local discriminator, (c) the result of using the multi-discriminator, and (d) the ground truth image. Best viewed with zoom-in.

Method	PSNR	SSIM	LPIPS
DSA without supervision	26.23	0.9062	0.0726
DSA with supervision	26.36	0.9107	0.0706

Table 4. Quantitative results over CelebA-HQ to prove the effect of supervision on DSA. Higher SSIM and PSNR values are better; lower LPIPS are better.

s in CA. Therefore, DSA can employ small-scale features as reference for inpainting. Furthermore, as DSA learns foreground self-attention, it maintains the contextual consistency in foreground areas. For example, the eyes generated by DSA are symmetric in both structure and texture, as shown in Figure 6. In comparison, the eyes generated by CA may be inconsistent with each other. Besides, DSA is also more efficient than CA as the latter one adopts more convolution operations [32]. The CSA module [15] also learns correlations between foreground patches, but its computational cost is significantly higher than DSA. This is because it adopts an iterative processing strategy, which means foreground pixels are refined one by one. Under the same GPU and experimentation setting, DSA and CSA (our implementation) take 6.4 ms and 31.3 ms per image on layer 14, respectively.

Effect of Multi-discriminator As shown in Figure 7, comparisons in visual effect show that using global and local discriminator only [13, 6] is insufficient to generate high-fidelity facial textures. With the help of local subdivision discriminator and four discriminators for facial components, the generated results are sharper and contain richer textures especially for the facial components.

Effect of Supervision on Attention As shown in Figure 8 (b), DSA without the extra supervision signal may produce artifacts for some challenging cases. For example, redundant hair is generated on the cheek of the girl in the first row. In the second row, the generated eye is close to




Figure 8. Qualitative comparison results on CelebA-HQ. In each row from left to right: (a) input, (b) the result of using DSA module but without supervision signal, (c) the result of using DSA module with supervision signal and (d) the ground truth image. The yellow circles indicate artifacts. Best viewed with zoom-in.

blue due to the influence of the blue background. In addition, there is clutter on the boundary between hair and background. These artifacts are caused by inaccurate attention scores, which lead the network to refer to features in unsuitable areas for inpainting. As illustrated in Figure 8 (c), after imposing the oracle supervision signal as a guidance, the accuracy of attention scores is improved and accordingly these artifacts or color discrepancy can be solved. As noted in Table 4, the quantitative results also prove that the network benefits from the supervision signal.

5. Discussion and Extension

In this paper, we propose a comprehensive model for face completion consisting of a DSA module with supervision signal, and a multi-discriminator, and further conduct multiple experiments to prove that it outperforms previous state-of-the-art methods. The supervised DSA module helps the network to identify the correlation between different facial parts, while the multi-discriminator forces the network to learn the specific features of facial components. In the supplementary material, we further show the results of the proposed method conditioning on irregular masks and higher resolution face image (1024×1024).

6. Acknowledgement

Changxing Ding is the corresponding author and he was supported by NSF of China under Grant 61702193 and U1801262, the Science and Technology Program of Guangzhou under Grant 201804010272, the Program for Guangdong Introducing Innovative and Entrepreneurial Teams under Grant 2017ZT07X183, and the Fundamental Research Funds for the Central Universities of China under Grant 2019JQ01.

References

- [1] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. In *ToG*, volume 28, page 24. ACM, 2009.
- [2] Antonio Criminisi, Patrick Pérez, and Kentaro Toyama. Region filling and object removal by exemplar-based image inpainting. *TIP*, 13(9):1200–1212, 2004.
- [3] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, pages 3146–3154, 2019.
- [4] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *NeurIPS*, pages 5767–5777, 2017.
- [5] Kaiming He and Jian Sun. Statistics of patch offsets for image completion. In *ECCV*, pages 16–29. Springer, 2012.
- [6] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):107, 2017.
- [7] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, pages 1125–1134, 2017.
- [8] 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.
- [9] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019.
- [10] Vahid Kazemi and Josephine Sullivan. One millisecond face alignment with an ensemble of regression trees. In *CVPR*, pages 1867–1874, 2014.
- [11] Jino Lee, Dong-Kyu Lee, and Rae-Hong Park. Robust exemplar-based inpainting algorithm using region segmentation. *TCE*, 58(2):553–561, 2012.
- [12] Ang Li, Jianzhong Qi, Rui Zhang, and Ramamohanarao Kotagiri. Boosted gan with semantically interpretable information for image inpainting. In *IJCNN*, pages 1–8. IEEE, 2019.
- [13] Yijun Li, Sifei Liu, Jimei Yang, and Ming-Hsuan Yang. Generative face completion. In *CVPR*, pages 3911–3919, 2017.
- [14] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *ECCV*, pages 85–100, 2018.
- [15] Hongyu Liu, Bin Jiang, Yi Xiao, and Chao Yang. Coherent semantic attention for image inpainting. In *ICCV*, 2019.
- [16] Ivan Molodetskikh, Mikhail Erofeev, and Dmitry Vatolin. Perceptually motivated method for image inpainting comparison. *arXiv preprint arXiv:1907.06296*, 2019.
- [17] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212*, 2019.
- [18] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, pages 2536–2544, 2016.
- [19] Yurui Ren, Xiaoming Yu, Ruonan Zhang, Thomas H Li, Shan Liu, and Ge Li. Structureflow: Image inpainting via structure-aware appearance flow. In *ICCV*, pages 181–190, 2019.
- [20] Tijana Ružić and Aleksandra Pižurica. Context-aware patch-based image inpainting using markov random field modeling. *TIP*, 24(1):444–456, 2014.
- [21] Min-cheol Sagong, Yong-goo Shin, Seung-wook Kim, Seung Park, and Sung-jea Ko. Pepsi: Fast image inpainting with parallel decoding network. In *CVPR*, pages 11360–11368, 2019.
- [22] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *Computer Science*, 2014.
- [23] Linsen Song, Jie Cao, Linxiao Song, Yibo Hu, and Ran He. Geometry-aware face completion and editing. In *AAAI*, pages 2506–2513, 2019.
- [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 *ECCV*, September 2018.
- [25] Yuhang Song, Chao Yang, Yeji Shen, Peng Wang, Qin Huang, and C-C Jay Kuo. Spg-net: Segmentation prediction and guidance network for image inpainting. In *BMVC*, page 97, 2018.
- [26] Yi Wang, Xin Tao, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. Image inpainting via generative multi-column convolutional neural networks. In *NeurIPS*, pages 331–340, 2018.
- [27] Chaohao Xie, Shaohui Liu, Chao Li, Ming-Ming Cheng, Wangmeng Zuo, Xiao Liu, Shilei Wen, and Errui Ding. Image inpainting with learnable bidirectional attention maps. In *ICCV*, pages 8858–8867, 2019.
- [28] Wei Xiong, Jiahui Yu, Zhe Lin, Jimei Yang, Xin Lu, Connelly Barnes, and Jiebo Luo. Foreground-aware image inpainting. In *CVPR*, pages 5840–5848, 2019.
- [29] Zhaoyi Yan, Xiaoming Li, Mu Li, Wangmeng Zuo, and Shiguang Shan. Shift-net: Image inpainting via deep feature rearrangement. In *ECCV*, pages 1–17, 2018.
- [30] Chao Yang, Xin Lu, Zhe Lin, Eli Shechtman, Oliver Wang, and Hao Li. High-resolution image inpainting using multi-scale neural patch synthesis. In *CVPR*, pages 6721–6729, 2017.
- [31] Raymond A Yeh, Chen Chen, Teck Yian Lim, Alexander G Schwing, Mark Hasegawa-Johnson, and Minh N Do. Semantic image inpainting with deep generative models. In *CVPR*, pages 5485–5493, 2017.
- [32] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *CVPR*, pages 5505–5514, 2018.
- [33] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *CVPR*, 2019.
- [34] Yanhong Zeng, Jianlong Fu, Hongyang Chao, and Baining Guo. Learning pyramid-context encoder network for high-quality image inpainting. In *CVPR*, pages 1486–1494, 2019.
- [35] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *ICML*, pages 7354–7363, 2019.

- [36] Haoran Zhang, Zhenzhen Hu, Changzhi Luo, Wangmeng Zuo, and Meng Wang. Semantic image inpainting with progressive generative networks. In *ACM-MM*, pages 1939–1947. ACM, 2018.
- [37] Hongying Zhang, Yuhong Jin, and Yadong Wu. Image completion by a fast and adaptive exemplar-based image inpainting. In *ICCAISM*, volume 3, pages V3–115. IEEE, 2010.
- [38] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018.
- [39] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. Pluralistic image completion. In *CVPR*, pages 1438–1447, 2019.