

Deep Fusion Network for Image Completion

Xin Hong

CAS Key Laboratory of Network Data Science and
Technology, Institute of Computing Technology, Chinese
Academy of Sciences
University of Chinese Academy of Sciences
Beijing, China
hongxin@ict.ac.cn

Pengfei Xiong, Renhe Ji, Haoqiang Fan

xiongpengfei@megvii.com
jirenhe@megvii.com
fhq@megvii.com
Megvii Technology
Beijing, China

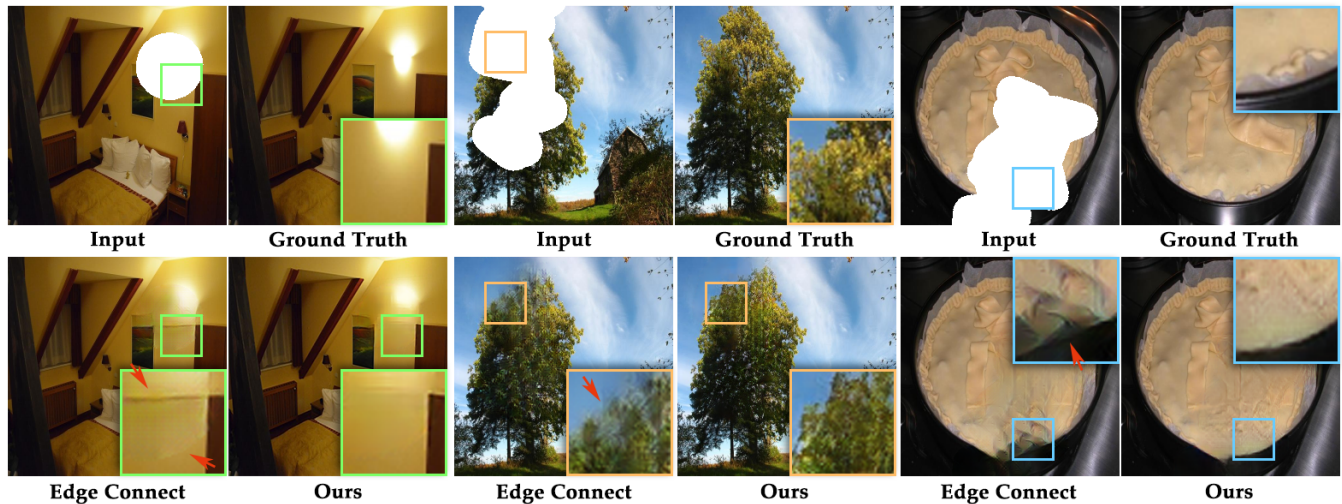


Figure 1: Comparison results between DFNet and previous state-of-the-art method Edge Connect[21]. In the first image of each group, white pixels represent the unknown region. With fusion blocks along with multi-scale constraints, DFNet has a smoother transition (1st case), more natural texture (2nd case) and more consistent structure (3rd case).

ABSTRACT

Deep image completion usually fails to harmonically blend the restored image into existing content, especially in the boundary area. This paper handles this problem from a new perspective of creating a smooth transition and proposes a concise Deep Fusion Network (DFNet). Firstly, a fusion block is introduced to generate a flexible alpha composition map for combining known and unknown regions. The fusion block not only provides a smooth fusion between restored and existing content but also provides an attention map to make network focus more on the unknown pixels. In this way, it builds a bridge for structural and texture information, so that information can be naturally propagated from the known region into completion. Furthermore, fusion blocks are embedded into several decoder layers of the network. Accompanied by the adjustable loss

constraints on each layer, more accurate structure information is achieved. We qualitatively and quantitatively compare our method with other state-of-the-art methods on Places2 and CelebA datasets. The results show the superior performance of DFNet, especially in the aspects of harmonious texture transition, texture detail and semantic structural consistency.

CCS CONCEPTS

• Computing methodologies → Image processing.

KEYWORDS

image completion, inpainting, Deep Fusion Network, Fusion Block, alpha composition

ACM Reference Format:

Xin Hong and Pengfei Xiong, Renhe Ji, Haoqiang Fan. 2019. Deep Fusion Network for Image Completion. In *Proceedings of the 27th ACM International Conference on Multimedia (MM '19)*, October 21–25, 2019, Nice, France. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3343031.3351002>

1 INTRODUCTION

Image completion, which aims to fill unknown regions of an image, is a fundamental task in computer vision. It can be broadly applied

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '19, October 21–25, 2019, Nice, France

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6889-6/19/10...\$15.00

<https://doi.org/10.1145/3343031.3351002>

to the fields of image editing, such as old photo recovering, object removal, and seamless inpainting for a damaged image. For most such applications, it is a critical problem to generate perceptually plausible completion results, specifically with natural transitions between known and unknown regions.

Previous approaches based on deep learning have shown great progress in image completion task [8, 10, 11, 18, 21, 24, 26, 30, 34–37]. As mentioned in [2], these methods can be divided into two groups. One group of works focus on building a contextual attention architecture or applying effective loss functions to generate more realistic content in the missing area. They assume the gaps should be filled with similar content from background. A typical arrangement is applying Partial Convolutions[18] to concentrate on the unknown region. Other methods regard structural consistency as the more important thing. Context priors such as edges are the most frequently used in these methods to ensure structural continuity. For instance, [21] proposed the Edge Connect method which can recover images with good semantic structural consistency. These approaches are dedicated to infer the unknown region with visually realistic and semantically related content. However, realizing a smooth transition is more critical than restoring texture-rich images in most scenarios, as shown in Figure 1.

Humans have an incredible ability to detect discontinuous transition regions. Consequently, The filled region must be perceptually plausible in the transition zone with sufficiently similar texture and consistent structure. To achieve smooth transitions, [25] proposed a method to iteratively optimize the pixel gradient in edge transitional region. Given two images, the fusion quality depends on the consistency of the gradient changes of these two images, which is similar to the relationship between the restored content and the known region in image completion. This inspires us to build a network to simulate the composition process.

In this work, we design a learnable fusion block to implement pixel-level fusion in the transition region. As shown in Figure 2, the fusion block can be embedded to an encoder-decoder structure. Different from the previous methods, we develop an extra convolutional block to generate an alpha map, which is similar to the hole mask but has smoother weights especially on the boundary region. In the process of gradient descent optimization, the alpha composition map adjusts the balance between restored image and ground truth content to make the transition smoother. Similar ideas have also been used in image matting [5, 22]. However, The purpose of these methods is to extract the smooth coefficients from background and foreground images, while the proposed fusion block is to combine them.

In detail, we propose a Deep Fusion Network (DFNet). Firstly, a fusion block is adopted as an adaptable module to combine the restored part of image and the original image. In addition to providing a smooth transition, the fusion block avoids learning unnecessary identity mapping for pixels in unknown regions and provides an attention map to make network focusing more on the missing pixels. With fusion block, structural and texture information can be naturally propagated from the known region into unknown regions. Secondly, we embed this module into different decoder layers. We find out that by considering the prediction of different fusion blocks with multi-scale constraints, the deep fusion network outperforms the network with only one fusion block embedded to the final layer.

Furthermore, while different layers produce different feature presentations, we selectively switch on and off structure and texture loss, to recover the structural information from lower layers and refine texture details in high layers. The whole architecture of DFNet is displayed in Figure 4.

The proposed DFNet is evaluated on two standard benchmarks, Places2 and CelebA. To better verify the proposed method, we define Boundary Pixels Error to measure the transition performance near the boundary of unknown regions. Also, ℓ_1 and FID are applied to verify global texture and consistency. Experiments demonstrate the superior performance of DFNet while compared with other state-of-the-art methods both in quantitative and qualitative aspects. It achieves better results in not only smooth texture transition but also structural consistency and more detailed textures.

In conclusion, the main contributions can be summarized as follows:

- We investigate the image completion problem with the perspective of better transition regions and propose fusion block which predicts an alpha composition map to achieve a smooth transition.
- Fusion block avoids learning unnecessary identity mapping for known regions and provides an attention mechanism. In this way, structure and texture information can propagate from known regions to completion more naturally.
- We propose Deep Fusion Network, a U-Net architecture embedded with multiple fusion blocks to apply multi-scale constraints.
- A new measurement Boundary Pixels Error (BPE) is introduced to measure the transition performance near the boundary of missing holes.
- The results on Places2 and CelebA show that our method outperforms state-of-the-art methods in both qualitative and quantitative aspects.

2 RELATED WORK

Context Aware Context-aware based image completion methods imagine the semantic content can be filled based on the overall scene. Context Encoders[24] introduces a encoder-decoder network to restore images from damaged inputs and holes. It applies a discriminator to increase the authenticity of restored images. Yang et al.[33] takes its result as input and then propagates the texture information from unknown regions to fill the missing area. Li et al.[17] and Iizuka et al.[11] extend Context Encoders by defining both global and local discriminators to pay more attention to the missing areas. Iizuka et al. applies Poisson Blending[25] as post-processing. Liu et al. [18] introduce partial convolution layers to avoid capturing too many zeros from unknown regions. These methods depend entirely on the training image to generate semantically relevant structures and texture confidence.

Texture Generation In the field of texture generation, a perceptual loss is adopted to fill in visually realistic content for missing regions. Liu et al.[18] applies perceptual loss[6, 14] which uses a VGG[29] network as a feature extractor. It computes loss use extracted high-level features to achieve higher resolution textures in completion. Other methods usually rely on GAN[7] loss to obtain better details. For instance, Yu et al.[36] replace the post-processing

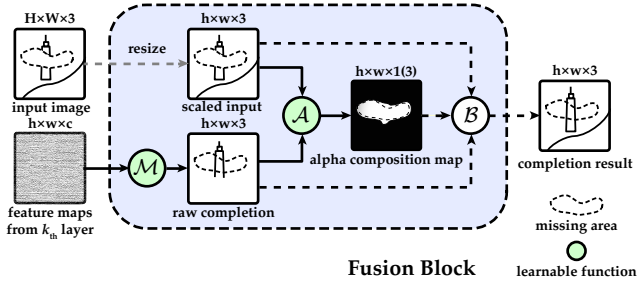


Figure 2: Illustration of Fusion Block. A fusion block extracts raw completion from feature maps by a learnable function \mathcal{M} , and also predicts an alpha composition map with the function \mathcal{A} . Finally, it combines the raw completion with the scaled input image by the blending function \mathcal{B} . The detail of blocks can be found in Section 3.1.

with a refinement network powered by the contextual attention layers.

Structure constraints To better control the completing behavior of networks, other works[21, 30, 35] explore providing extra information for inpainting. Song et al.[30] uses a DeepLabv3+[3] model to first predict a segmentation map, and then completes the unknown region with predicting segmentation map as prior. Yu et al.[35] propose gated convolution which generalizes partial convolution and the new structure is compatible with user guides, usually strokes to indicate edges. Like Song et al.[21] uses a two-staged networks for completion. It first completes edges corresponding to the input image and then use completed strokes to guide the completing of color images. To some extent, those methods can manually control the completion result of networks by replacing the priors with custom one or giving extra edge information.

Image Embedding As a similar work with image completion, image embedding and matting are also studied in the past decades. [25] proposes a method to iteratively optimize the pixel gradient in edge transitional region. Then Poisson matting[31] firstly introduces a Poisson blending method into alpha matting by solving a Poisson equation, which proves the effectiveness of alpha composition. Deep Image matting[22] also generates an alpha map with an encoder-decoder network. Cho et al[5] take the matting results of [1] and normalized RGB colors as inputs and learn an end-to-end deep network to predict a new alpha matte. These methods prove that alpha matting based on deep learning is more realistic for image embedding and matting.

3 DEEP FUSION NETWORK

Deep Fusion Network is built on a U-Net[27] like architecture, which is widely used in recent image segmentation[20] task and image to image translation[4, 13, 16, 32] tasks. The difference between our DFNet and original U-Net is that we embed fusion blocks to several layers of a decoder. Fusion blocks help us to achieve smoother transitions near the boundary and are the key components for our multi-scale constraints. In this section, we first introduce the fusion block and then discuss our network architecture and loss functions.




Figure 3: Corresponding results in a fusion block.

3.1 Fusion Block

The task of image completion is to restore the missing area with visually plausible content from a damaged image I_{in} and a binary mask M which represents the location of the unknown region.

Recently deep learning based methods usually predict the whole image I_{out} which even includes known regions and use it to calculate loss during training. However, they take I_{comp} ($I_{comp} = M \odot I_{in} + (1 - M) \odot I_{out}$, \odot denotes Hadamard product) rather than I_{out} for testing. The composition process replaces known region in I_{out} with corresponding pixels in I_{in} . Only a few methods[18] use both I_{out} and I_{comp} to compute loss.

This training strategy has problems. Firstly, the mission of image completion is to complete the unknown region only. However, it's hard to complete missing holes while keeping a strict identity mapping for known area. Secondly, the inconsistent use of I_{comp} and I_{out} during training and testing, along with the rigid composition method, usually produces visible artifacts around the boundary of missing areas. As shown in the first case of Figure 1, the result of Edge Connect[21] has a clear edge at the boundary of completion.

To remove the artifacts around the boundary and avoid neural networks learning unnecessary identity mapping, we propose *Fusion Block*. As shown in Figure 2, a fusion block feed with two elements, an input image with unknown region I_{in} and feature maps F_k from k_{th} layer (1st layer is the last decoder layer of U-Net). The fusion block first extracts raw completion C_k from feature maps, and then predicts an alpha composition map α_k to combine them. The final result \hat{I}_k is obtained by:

$$\hat{I}_k = \mathcal{B}(\alpha_k, C_k, I_k) = \alpha_k \odot C_k + (1 - \alpha_k) \odot I_k$$

We resize I_{in} to obtain I_k . The raw completion C_k extracted from feature maps F_k by a learnable function \mathcal{M} :

$$C_k = \mathcal{M}(F_k)$$

$\mathcal{M}(x)$ transforms n channel feature maps x into a 3 channel image with the resolution unchanged which is exactly the raw completion. In practice, we use a 1×1 convolutional layer following with a sigmoid function to learn \mathcal{M} .

The alpha composition map α_k is produced by another learnable function \mathcal{A} from raw completion and the scaled input image:

$$\alpha_k = \mathcal{A}(\mathbf{F}_k, \hat{\mathbf{I}}_k)$$

α_k has two choices in the number of channels, either single channel for image-wise alpha composition or 3 channels for channel-wise alpha composition. In practice, we find channel-wise alpha composition performs better. As for \mathcal{A} , we use three convolutional layers and the kernel size of them is 1, 3, 1. First two convolutional layers are followed with a Batch Normalization[12] layer and a leaky ReLU function. And we apply a sigmoid function to the output of the last convolutional layer.

A fusion block enables a network to avoid learning unnecessary identity mapping while completing unknown regions with soft transition near the boundary. We also give an example of corresponding images in a fusion block in Figure 3. Completion performance can be further improved with multi-scale constraints by embedding fusion blocks to the last few decoder layers of U-Net.

3.2 Network Architecture

It's intuitive that when completing an image, constructing structures is easier in a lower resolution for algorithms, while recovering texture is more feasible in a higher resolution. We embed fusion blocks to the last few decoder layers of the U-Net and obtain completion results in different resolutions. And then we can apply structure and texture constraints to different resolution as we want. The overview of our DFNet is shown in Figure 4. We choose U-Net[27] like the one used in [13, 18] as our backbone architecture. The difference is that the last few decoder layers are embedded with fusion blocks. Each fusion block outputs a completion result \mathbf{C}_i with the same resolution as the input feature maps \mathbf{F}_i . According to their resolution, we can provide different constraints as we want during training. We will discuss these constraints in Section 3.3. During testing, only the completion result $\hat{\mathbf{I}}_0$ from the last layer is needed.

3.3 Loss Functions

The target of image completion is to generate visually plausible results in both aspects of structure and texture. Reconstruction loss, which is mean absolute error of each pixel between prediction and ground truth, is usually used to guarantee accurate structures in completion results. However, high-resolution textures are beyond the capability of reconstruction loss. Previous works use GAN loss [7] or perceptual loss along with style loss [14] to obtain vivid textures. But they have some drawback which is known as producing checkerboard and fish scale artifacts[18]. Total variation loss is usually used to counter this drawback. Results from [18] shows that this artifact can be reduced more obviously by increasing the weight of style loss.

Reconstruction Loss Reconstruction loss is defined as mean absolute error of completion result $\hat{\mathbf{I}}_k$ and target image \mathbf{I}_k :

$$\mathcal{L}_{\ell_1}^k = \frac{1}{C_k H_k W_k} \|\mathbf{I}_k - \hat{\mathbf{I}}_k\|_1$$

The number of channels is C , the height is H and the width is W .

Perceptual Loss and Style Loss Perceptual loss and style loss are first used in style transfer[6, 14]. They use a pre-trained VGG network to extract high-level features. The errors are computed between these features rather than original images. Let $\phi_j(x)$ be the features of j th layer in a VGG network when given an image x . The size of $\phi_j(x)$ is $C_j \times H_j \times W_j$. Perceptual loss is defined as the error of these features:

$$\mathcal{L}_p^k = \sum_{j \in J} \|\phi_j(\mathbf{I}_k) - \phi_j(\hat{\mathbf{I}}_k)\|_1$$

J is selected VGG layers. *Gram matrix* is a $C_j \times C_j$ matrix, whose elements are defined as:

$$G_j^\phi(x)_{c,c'} = \frac{1}{C_j H_j W_j} \sum_{h=1}^{H_j} \sum_{w=1}^{W_j} \phi_j(x)_{h,w,c} \phi_j(x)_{h,w,c'}$$

And then style loss is L_1 mean absolute error between corresponding Gram matrices of the output and target image:

$$\mathcal{L}_s^k = \sum_{j \in J} \|G_j^\phi(\mathbf{I}_k) - G_j^\phi(\hat{\mathbf{I}}_k)\|_1$$

Style loss doesn't consider the position of pixels but cares about how high level features appear simultaneously[14], so that it's better for constraining the entire style of an image.

Total Variation Loss Total variation loss \mathcal{L}_{tv} is errors computed only use predictions. Each pixel will compute errors with top pixel and left pixel respectively. This can be implemented more easily by using a convolution layer with a fixed kernel.

Total Loss. We group loss functions into *Structure Loss* and *Texture Loss*. Structure Loss is represented as weighted reconstruction loss:

$$\mathcal{L}_{struct}^k = \lambda_{\ell_1} \mathcal{L}_{\ell_1}^k$$

And texture loss is a combination of three loss:

$$\mathcal{L}_{text}^k = \lambda_p \mathcal{L}_p^k + \lambda_s \mathcal{L}_s^k + \lambda_{tv} \mathcal{L}_{tv}^k$$

Our final loss is sum of structure loss and texture loss from different resolution completion results:

$$\mathcal{L}_{total} = \frac{1}{|P|} \sum_{k \in P} \mathcal{L}_{struct}^k + \frac{1}{|Q|} \sum_{k \in Q} \mathcal{L}_{text}^k$$

P is the set of layers which consider structure loss while Q is for texture loss. And for brevity, we use (p, q) to represent the choice of P, Q , which takes last p layers as P and last q layers as Q . For example, $(2, 1)$ represent $P = \{1, 2\}$ and $Q = \{1\}$, which means completion results from the last two layers of U-Net will be used to compute structure loss and only last one for texture loss. The corresponding part will be ignored if the total number of layers $|P|$ or $|Q|$ equal to zero. We will discuss the choice of P and Q in Section 4.3.

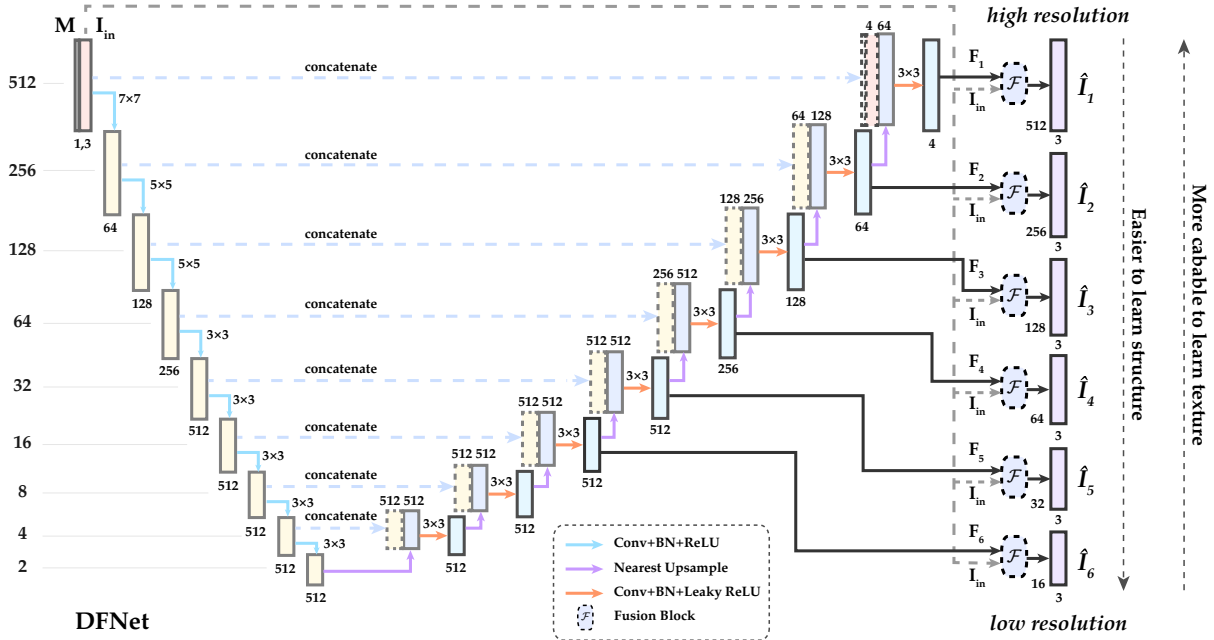


Figure 4: Overview of our *Deep Fusion Network (DFNet)*. DFNet is based on a U-Net, like the one used in [13, 18]. The difference from traditional U-Net is that we embed fusion blocks to the last few decoder layers. During training, each fusion block will produce a completion result \hat{I}_k from corresponding feature maps, which also has the same resolution with the feature maps. So that different constraints can be provided to each completion result as needed. During testing, only the completion results from the last layer need to be produced.

4 EXPERIMENTS

4.1 Experiment Details

We evaluate DFNet on two public datasets: **Places2**[38] and **CelebA** [19]. For Places2, we use the original train, test, and val splits. For CelebA, we randomly partition into 27K images for training and 3K images for testing. Images in Places2 and CelebA are respectively resized to 512×512 and 256×256 during training and testing. We randomly generate 1000 masks according to the method in [35] and perform augmentation to these masks during training. To analysis the influence of unknown region range, these masks are categorized into five classes, including [0, 10%), [10%, 20%), [20%, 30%), [30%, 40%), [40%, 50%).

Models are separately trained on each dataset. Our proposed model is implemented in PyTorch[23] and trained in a single machine with 8 GeForce RTX 2080 Ti. We use Horovod[28] as our distributed training framework. With a batch size of 6 for each GPU, it usually takes about 3 days to train a model. Forwarding is extremely fast, it only takes 8.29ms to complete an image. As a common configuration, Adam[15] is applied for optimization. The learning rate reduced from $2e-3$ to $2e-6$ in 20 epochs, with a decay rate of 0.1 and a step size of 5.

4.2 Evaluation Metrics

Different from the tasks as image classification, detection, and segmentation, image generation usually doesn't have strict targets. The

basic rule is visually plausible. For image completion, it requires the completion not only looks real but also transit naturally from known regions. So we apply ℓ_1 , and *Fréchet Inception Distance* (FID) [9] as evaluation metrics both in perspective of pixels and features to quantitatively analyze the performance of DFNet.

Furthermore, we observe that pixels in unknown region that near the boundary have very small variance while these pixels play the most important role in structure and texture transition. To measure the transition performance of models, we propose *Boundary Pixels Error* (BPE) which only consider pixels error near the boundary. For boundary area b , which is n pixels narrow band adjacent to the boundary of unknown region, BPE is mean absolute error of those pixels between ground truth I and prediction \hat{I} :

$$BPE = \frac{\|b \odot (I - \hat{I})\|_1}{\|b\|_1}$$

4.3 Analysis of DFNet architecture

In this section, we investigate the performance of the proposed modules in DFNet. First, we show the effectiveness of fusion blocks. Then we focus on the effect of multi-scale constraints by gradually increasing fusion blocks to DFNet and evaluating it. Finally, we discuss how to apply structure loss and texture loss on different resolution completion results to achieve the best results.

4.3.1 Effectiveness of Fusion Block. We compare our results with predictions from a normal U-Net and predictions from a DFNet

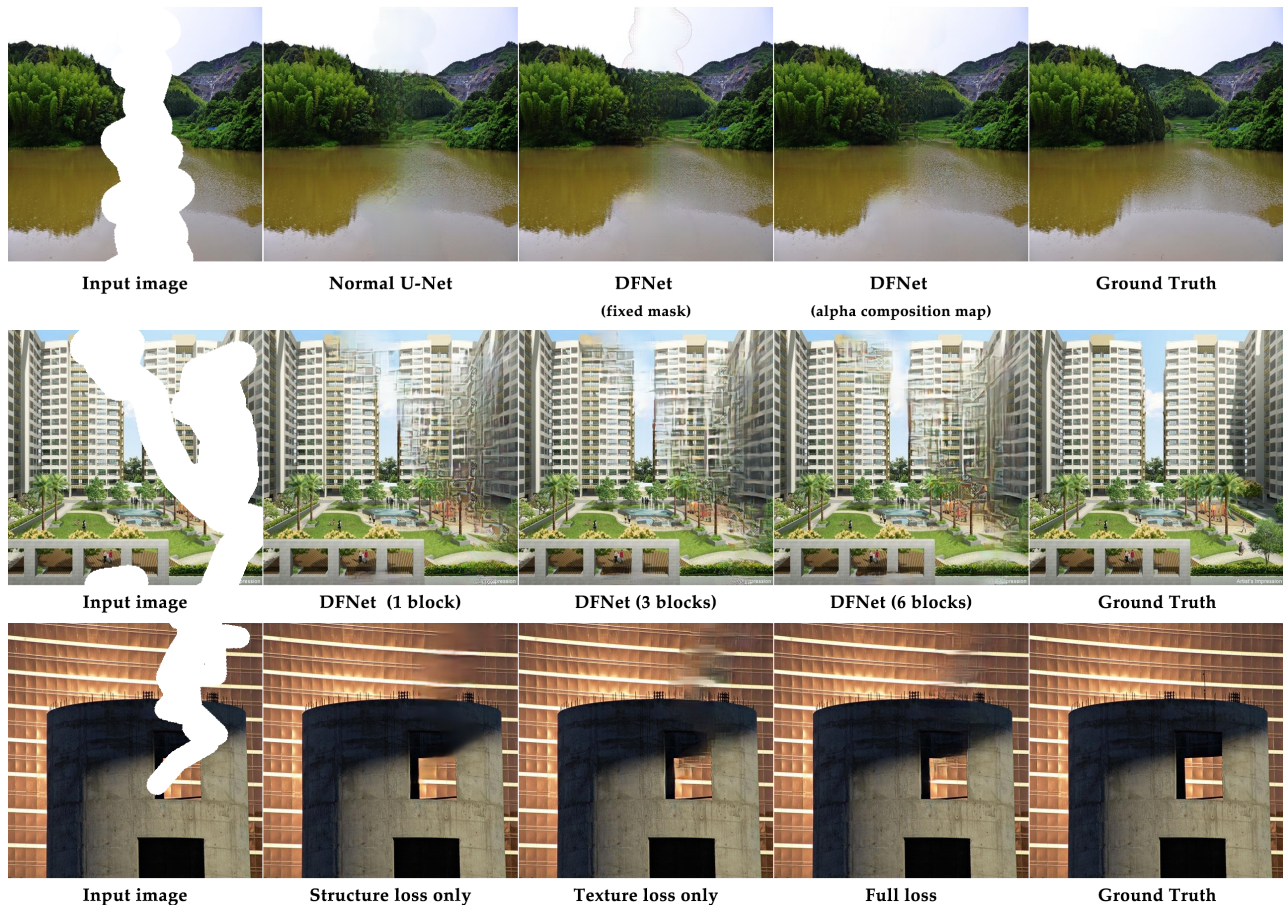


Figure 5: Effectiveness of Fusion Block, Multi-scale constraints, and Loss Ablation. (a) compares the results of the proposed network without and with a fixed mask, and with a fusion block. (b) depicts the results with 1, 3, 6 fusion blocks respectively. (c) shows the effects of structure loss and texture loss, and proves the effectiveness of combination loss.

but directly using the mask to replace the alpha composition map. We use only one fusion block for fair comparison, which means $P = Q = \{1\}$ for \mathcal{L}_{total} . (Section 3.3).

As can be seen in the 1st row of Figure 5, fusion blocks lead to the best transitions near the boundary. Although most of the semantic information has been restored, there exists obvious color transition inconsistent in the result of the standard network without mask constraints. This is because global semantic consistency constraints can only lead to similar texture in the missing areas, but structural consistency can not be guaranteed. Based on the mask constraints, the pixel transition in filling areas becomes more natural, which proves the effect of the proposed method on the propagation of structural and texture information. As mentioned above, the alpha composition map is an attention mechanism to enhance structural consistency. Furthermore, the result of the learned alpha mapping is even better in the edge transition to eliminate the visible artifacts near the boundary.

The same detailed conclusion can be seen in Figure 3. Based on the proposed fusion block, the structure between the known and unknown areas are well preserved, even beyond the mask area. The

sharp edge of the roof is retained into the reconstructed image with other useless parts discarded.

4.3.2 Multi-scale constraints. We compare DFNet with different numbers of fusion blocks from one to six. Formally speaking, P and Q in Section 3.3 both increase from $\{1\}$ to $\{6\}$. In this section, P and Q are equal to only analyze the role of multi-scale fusion.

As can be seen in the 2nd row of Figure 5, the structure of the building is more clear and accurate based on more fusion blocks. Also, the shapes of houses are depicted in the result of 3 fusion blocks instead of the noises in the result based on 1 fusion block. While high-level layers in encoder have a bigger receptive field and global context, the structure information can be more easily reconstructed with more layers in the decoder. Nevertheless, although the result of 6 fusion blocks retains this structural information, its texture is not very stable compared to 3 fusion blocks. We guess this is because we shouldn't apply texture constraints for low-resolution completion results. In the next section, we will go into more detail about how to choose the number of blocks layers.

We also give the quantitative comparison in the second row of Figure 6. Results are separated according to the range of mask




Figure 6: Evaluation of different P and Q . Detailed description for P and Q can be found in Section 3.3. We separately compare models with different area range of mask. The results have been normalized by subtracting the minimal value of corresponding comparisons. For these three metrics, lower means better performance. In (a), (b) and (c), we compare 6 choices of P and Q which gradually add fusion blocks but keeping $P = Q$. Results show performance gets better with multi-scale constraints. We further compare other 4 choices of P, Q which $P \neq Q$ and choose (6, 3) as our best model. More analysis can be found in Section 4.3.2 and Section 4.3.3.

in each evaluation metric. With fusion blocks increased, FID gets lower and lower. This means multi-layer constraints help to capture contextual information and make the whole image looks more real. The BPE increases slightly with fusion blocks increased. This can be explained that a finer texture and a smoother transition is a trade-off. However, the globally visual effect is more important while the gaps on BPE is relatively small.

4.3.3 Loss Ablation and Tuning. Firstly, the effect of structure loss \mathcal{L}_{struct} and texture loss are showed \mathcal{L}_{text} by respectively trained DFNet only applies only one of them. As seen in the 3rd Figure 5, the result without texture loss is blurry but with accurately structure consistency, while the other one destroys the structure, it fails to recover edges of the object although they have finer textures. This provides strong evidence for loss design in this paper.

We further discuss the dynamic loss design in each layer. Based on the visualization results in 4.3.2, we make a comprehensive comparison of the loss design in different layers. As shown in Figure 6, the performance depicts the same trend with different ranges of hole size. We choose $P = \{1, 2, 3, 4, 5, 6\}$ to compute structure loss and $Q = \{1, 2, 3\}$ for texture loss in the final architecture. This can be explained that, although the structure information is more and more abundant with higher and higher encoder layers, the high-level features will lead to texture noise due to the loss of global semantic constraints.

4.4 Comparisons with Other Methods

We quantitatively and qualitatively compare our DFNet with 3 recently methods, including DeepFill [36], PConv [18] and Edge Connect[21]. Results of DeepFill and Edge Connect are obtained by using their pre-trained models^{1 2}. However, we don't find the official implementation of PConv, so we implement one with the same settings described in the original paper.

4.4.1 Quantitative Comparisons. Table 1 shows the comparison results on Places2[38]. We use three metrics including ℓ_1 , BPE and Fréchet Inception Distance (FID) [9]. Results from ours outperform others on both boundary transition and realistic on overall images. Our predictions on BPE is significantly lower than those from Edge Connect[21] and other methods. This means completion from our methods have better transitional areas near the boundary, which also proves the effectiveness of proposed fusion blocks.

Edge Connect works well on maintaining structural consistency by applying additional edge constraints. However, it doesn't pay much attention to the smooth transition. The constraints on the structure of the whole image can't lead to natural image restoration, especially in detail. Results of Edge connect shows lower ℓ_1 than ours while the missing hole is large. But this only state results of Edge Connect is more similar to original images. Because completion can be more diverse while the hole is larger.

¹https://github.com/JiahuiYu/generative_inpainting

²<https://github.com/knazeri/edge-connect/>

Table 1: Quantitative comparison with other methods on Places2.

Methods	ℓ_1					BPE					FID				
	0.1	0.2	0.3	0.4	0.5	0.1	0.2	0.3	0.4	0.5	0.1	0.2	0.3	0.4	0.5
DeepFill	2.79	6.75	10.63	15.35	28.38	1.33	1.81	2.39	2.91	5.13	24.04	56.55	98.25	173.90	324.97
PConv	1.51	4.22	7.01	10.52	12.83	0.17	0.37	0.62	0.87	1.60	14.98	41.21	84.60	166.72	217.48
EdgeConnect	1.43	3.94	6.41	9.64	11.38	0.33	0.69	1.11	1.48	2.55	19.24	35.91	68.29	131.16	147.51
DFNet	1.40	3.91	6.50	9.89	11.96	0.15	0.33	0.55	0.74	1.42	12.27	34.64	65.25	127.58	136.22

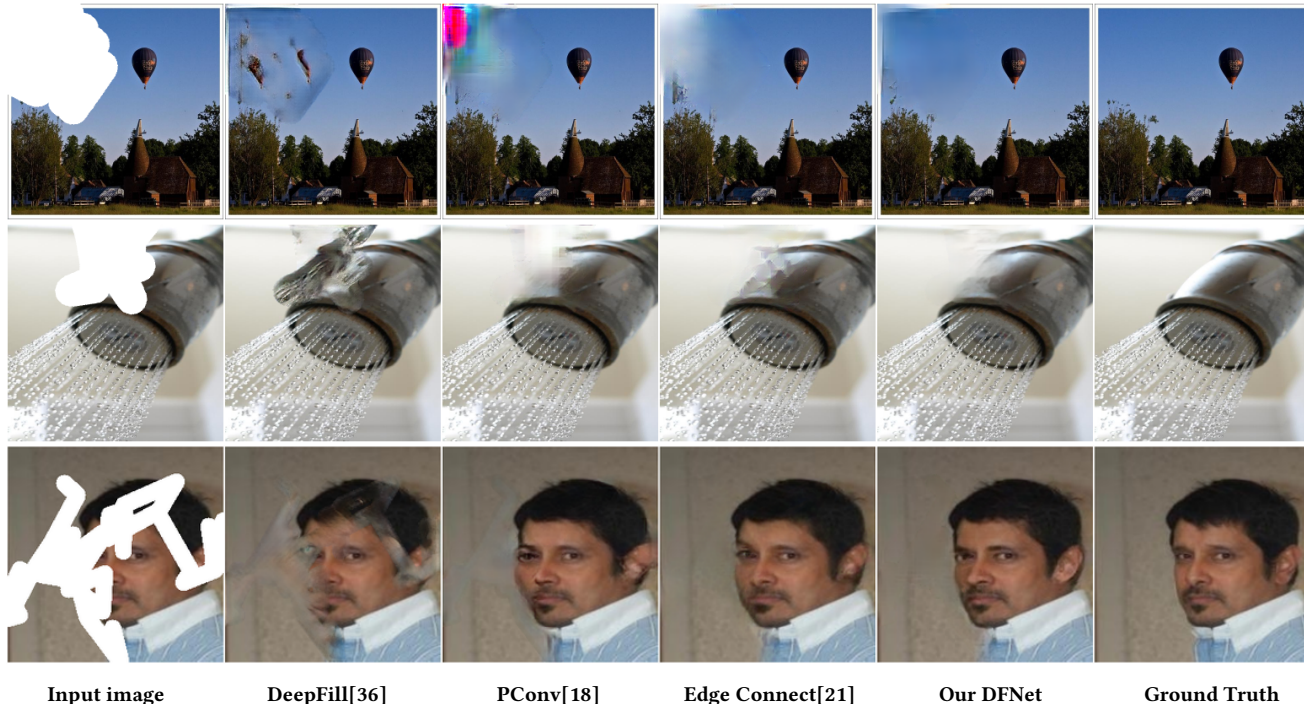


Figure 7: Comparison results on Places2 and CelebA. More results can be found in the supplementary materials.

PConv use partial convolution to progressively reduce missing region, which can be considered as providing a hidden attention map gradually enlarged from boundary area to full known region. This enhances the learning ability near the boundary, which have similar effects with the proposed DFNet when considering transition performance. However, this architecture is not good at a large hole because information can't be transmitted effectively to the inner area. When comparing PConv and Edge Connect on BPE and FID, we can find PConv has better transition near the boundary than Edge Connect and comparable FID when missing holes are small, however, when missing hole becomes larger, Edge Connect will have more realistic results.

4.4.2 *Qualitative Comparisons.* Figure 7 shows the comparison on Places2 and CelebA without any post-processing. As shown in the figure, we can see our model has the best performance in texture consistency near the boundary, and also good at keeping the structure consistency even better than Edge Connect. Results from different datasets show the generalization ability of our methods.

There is one thing that should be noticed, as shown in the 1st case of Figure 7, we find PConv and Edge Connect sometimes fail to complete the missing hole when the missing hole cover the border of an image. For PConv, we think this is the limit of partial convolution, which can't transmit information into a very large hole. While for Edge Connect, it always produces clouds like completions in a similar situation. We couldn't figure out the reason.

5 CONCLUSION

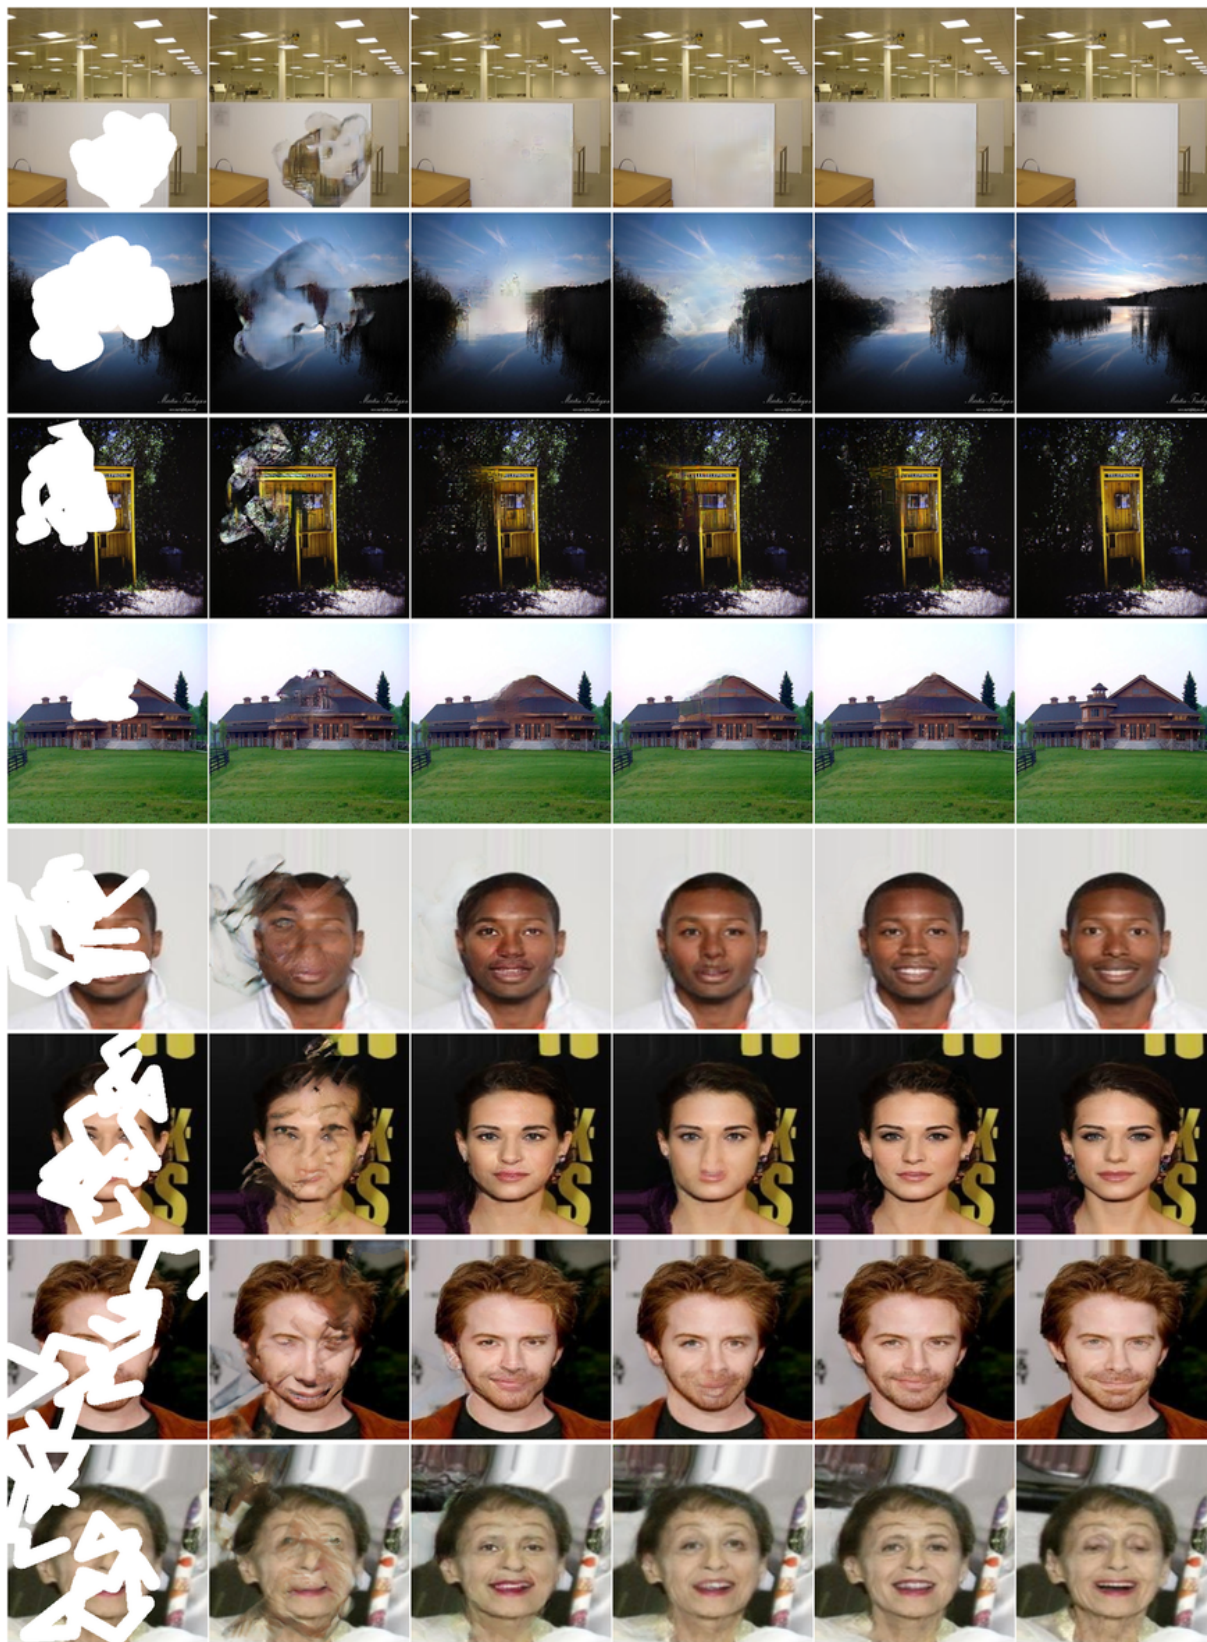
In this paper, we analyze the image completion technology from a new perspective. We propose Deep Fusion Network by designing a fusion block to predict an alpha composition map for combining completion and existing content and embedding it on multi-scale layers. Results of experiments on Places2 and CelebA dataset shows our method achieves state-of-the-art results, especially in the field of harmonious texture transition, texture detail and semantic structural consistency.

REFERENCES

- [1] D. Lischinski A. Levin and Y. Weiss. 2008. A closed-form solution to natural image matting. In *Pattern Analysis and Machine Intelligence*. 228–242.
- [2] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. 2000. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 417–424.
- [3] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 801–818.
- [4] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [5] I. Kweon D. Cho Y.-W. Tai D. Cho, Y.-W. Tai and I. Kweon. . 2016. Natural image matting using deep convolutional neural networks. In *European Conference on Computer Vision*. Springer, 1–8.
- [6] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2015. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576* (2015).
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [8] Changzhi Luo Wangmeng Zuo Meng Wang Haoran Zhang, Zhenzhen Hu. 2018. Semantic Image Inpainting with Progressive Generative Networks. In *ACM Multimedia*.
- [9] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*. 6626–6637.
- [10] Patrick Perez Huy V. Vo, Ngoc Q. K. Duong. 2018. Structural inpainting. In *ACM Multimedia*.
- [11] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. 2017. Globally and Locally Consistent Image Completion. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2017)* 36, 4, Article 107 (2017), 107:1–107:14 pages.
- [12] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Francis Bach and David Blei (Eds.), Vol. 37. PMLR, Lille, France, 448–456. <http://proceedings.mlr.press/v37/loff15.html>
- [13] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. *CVPR* (2017).
- [14] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*.
- [15] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [16] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiří Matas. 2018. DeblurGAN: Blind Motion Deblurring Using Conditional Adversarial Networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [17] Yijun Li, Sifei Liu, Jimei Yang, and Ming-Hsuan Yang. 2017. Generative Face Completion. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [18] Guilin Liu, Fitsum A. Reda, Kevin J. Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. 2018. Image Inpainting for Irregular Holes Using Partial Convolutions. In *The European Conference on Computer Vision (ECCV)*.
- [19] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- [20] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully Convolutional Networks for Semantic Segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [21] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. 2019. EdgeConnect: Generative Image Inpainting with Adversarial Edge Learning. *arXiv preprint*.
- [22] Scott Cohen Thomas Huang Ning Xu, Brian Price. 2017. Deep Image Matting. *arXiv preprint arXiv:1703.03872* (2017).
- [23] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *NIPS-W*.
- [24] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei Efros. 2016. Context Encoders: Feature Learning by Inpainting. In *Computer Vision and Pattern Recognition (CVPR)*.
- [25] Patrick Pérez, Michel Gangnet, and Andrew Blake. 2003. Poisson image editing. *ACM Transactions on graphics (TOG)* 22, 3 (2003), 313–318.
- [26] Teck Yian Lim Alexander G. Schwing Mark Hasegawa-Johnson Minh N. Do Raymond A. Yeh, Chen Chen. 2018. Semantic Image Inpainting With Deep Generative Models. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [27] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.
- [28] Alexander Sergeev and Mike Del Balso. 2018. Horovod: fast and easy distributed deep learning in TensorFlow. *arXiv:cs.LG/1802.05799*
- [29] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [30] Yuhang Song, Chao Yang, Yeji Shen, Peng Wang, Qin Huang, and C-C Jay Kuo. 2018. Spg-net: Segmentation prediction and guidance network for image inpainting. *arXiv preprint arXiv:1805.03356* (2018).
- [31] Jian Sun, Jiaya Jia, Chi-Keung Tang, and Heung-Yeung Shum. 2004. Poisson matting. In *ACM Transactions on Graphics (ToG)*, Vol. 23. ACM, 315–321.
- [32] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [33] Chao Yang, Xin Lu, Zhe Lin, Eli Shechtman, Oliver Wang, and Hao Li. 2017. High-Resolution Image Inpainting Using Multi-Scale Neural Patch Synthesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [34] Xiaojuan Qi Xiaoyong Shen Jiaya Jia Yi Wang, Xin Tao. 2018. Image Inpainting via Generative Multi-column Convolutional Neural Networks. In *Conference and Workshop on Neural Information Processing Systems(NeurIPS)*.
- [35] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. 2018. 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. 2018. Generative Image Inpainting with Contextual Attention. *arXiv preprint arXiv:1801.07892* (2018).
- [37] Zhe Lin Xiaofeng Liu Qin Huang Hao Li C.-C. Jay Kuo Yuhang Song, Chao Yang. 2018. Contextual-based Image Inpainting: Infer, Match, and Translate. In *The European Conference on Computer Vision (ECCV)*.
- [38] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. Places: A 10 million Image Database for Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017).

A MORE RESULTS

We provide more comparison results on Places2 and CelebA in this section.



Input image

DeepFill

PConv

Edge Connect

Our DFNet

Ground Truth

Figure 8: comparison results on CelebA.