

# Deep Portrait Image Completion and Extrapolation

Xian Wu<sup>1</sup>, Rui-Long Li, Fang-Lue Zhang<sup>2</sup>, *Member, IEEE*, Jian-Cheng Liu, Jue Wang<sup>3</sup>, *Senior Member, IEEE*,  
Ariel Shamir, and Shi-Min Hu<sup>4</sup>, *Senior Member, IEEE*

**Abstract**—General image completion and extrapolation methods often fail on portrait images where parts of the human body need to be recovered - a task that requires accurate human body structure and appearance synthesis. We present a two-stage deep learning framework for tackling this problem. In the first stage, given a portrait image with an incomplete human body, we extract a complete, coherent human body structure through a human parsing network, which focuses on structure recovery inside the unknown region with the help of full-body pose estimation. In the second stage, we use an image completion network to fill the unknown region, guided by the structure map recovered in the first stage. For realistic synthesis the completion network is trained with both perceptual loss and conditional adversarial loss. We further propose a face refinement network to improve the fidelity of the synthesized face region. We evaluate our method on publicly-available portrait image datasets, and show that it outperforms other state-of-the-art general image completion methods. Our method enables new portrait image editing applications such as occlusion removal and portrait extrapolation. We further show that the proposed general learning framework can be applied to other types of images, e.g. animal images.

**Index Terms**—Image completion, portrait extrapolation, human parsing, deep learning.

## I. INTRODUCTION

THERE are common mistakes that novice users often make when shooting a portrait photo. As a typical example, while photography rules suggest that cutting off hands, feet, and foreheads can ruin the visual flow, many portrait images are taken with such improper composition. The feet are often cut off as the photographer is focusing mostly on the face region when shooting the picture (see Fig. 1b). At other times, the accessories that the person is carrying (e.g. the bag in Fig. 1a), or the other objects that partially occlude the main subject (e.g. the dog in Fig. 12) could be distracting

and better be removed from the photo. To remove these imperfections, one could specify the unwanted objects to remove, or expand the border of the image to try to cover the whole human body, both leaving holes or blank regions in the image to be filled in properly.

Although image completion has been actively studied in the last twenty years, there is no existing approach that can work well for portrait images, where holes on the human body need to be filled in. A successful completion method is required to recover not only the correct semantic structure in the missing region, but also the accurate appearance of the missing object parts. Both goals are challenging to achieve for portrait photos. In terms of semantic structures, although human body has a very constrained 3D topology, the large variation of its pose configurations, coupled with 3D to 2D projection, makes accurate structure estimation and recovery from a single 2D image difficult. In term of appearance, despite the large variation in clothing, there are strong semantic constraints such as symmetry that the completion algorithm has to obey, e.g. two shoes usually have the same appearance regardless of how far they are separated in the image. As we will show later, without paying special attention to these constraints, general-purpose hole filling methods often fail to generate satisfactory results on portrait photos (see Fig. 7 and 8).

Extracting human body structures from images and videos has been well studied in previous human parsing [1], [2] and pose estimation approaches [3], [4]. Naturally, we want to rely on these methods to estimate human body structure from an incomplete portrait photo first, and then use it to guide the image completion process. Following this idea, we propose a two-stage deep learning framework for portrait photo completion and expansion. In the first stage, we utilize a human parsing network to estimate both of a human pose and a parsing map simultaneously from the input image. The human pose is then used to help refine the parsing map, especially inside the unknown region. In the second stage, we employ an image completion network to synthesize the missing region in the input image with the guidance of the parsing map, followed by a face refinement network for improving the generated face region. These two networks are trained sequentially with both the perceptual loss and the adversarial loss for improved realism. We show that our approach can be used in many portrait image editing tasks that are difficult or impossible for traditional methods to achieve, such as portrait extrapolation for recovering missing human body parts (e.g. Fig. 11), and occlusion removal (e.g. Fig. 12). Furthermore, We demonstrate that the proposed learning framework is quite generic, and is

Manuscript received August 19, 2018; revised May 5, 2019 and August 14, 2019; accepted October 1, 2019. Date of publication October 11, 2019; date of current version January 10, 2020. This work was supported by the National Natural Science Foundation of China under Project 61561146393 and Project 61521002. The work of F.-L. Zhang was supported by the Research Establishment Grant of the Victoria University of Wellington under Project 8-1620-216786-3744. The work of A. Shamir was supported by the Israel Science Foundation under Project 2216/15. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Ce Zhu. (*Corresponding author: Shi-Min Hu.*)

X. Wu, R.-L. Li, J.-C. Liu, and S.-M. Hu are with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China (e-mail: shimin@tsinghua.edu.cn).

F.-L. Zhang is with the School of Engineering and Computer Science, Victoria University of Wellington, Wellington 6012, New Zealand.

J. Wang is with Megvii (Face++) Research, Seattle, WA 98052 USA.

A. Shamir is with the Efi Arazi School of Computer Science, Interdisciplinary Center Herzliya, Herzliya 4610101, Israel.

Digital Object Identifier 10.1109/TIP.2019.2945866

applicable for other types of images such as ones of horses and cows (e.g. Fig. 10).

To the best of our knowledge, we are the first to integrate deep human body analysis techniques into an image completion framework. Our main contributions of this paper are summarized as follows:

- we propose a novel two-stage deep learning framework where human body structure is explicitly recovered to guide human image completion with face refinement;
- we show that our framework enables new portrait photo editing capabilities such as re-composition by extrapolation.

## II. RELATED WORK

### A. Image Completion

Traditional image completion methods can be categorized into diffusion-based and patch-synthesis-based ones. Diffusion-based methods, first proposed by Bertalmio *et al.* [5] and then extended by Ballester *et al.* [6] and Bertalmio *et al.* [7], propagate nearby image structures to fill in the holes by assuming local continuity. These techniques however can only handle small holes or narrow gaps. Patch-based methods are derived from texture synthesis algorithms [8]–[10]. They extract patches from the known region of the image and use them to fill large holes. Criminisi *et al.* [11] proposed a best-first algorithm by searching the most similar patch. Simakov *et al.* [12] presented a global optimization approach based on bidirectional similarity. These techniques were greatly accelerated by PatchMatch [13], [14], a randomized nearest neighbor field algorithm. They were further improved with appending gradients into the distance metric by Darabi *et al.* [15]. These methods can handle larger holes than propagation-based ones, but still need semantic guidance for structured scenes.

Many inpainting approaches rely on additional guidance for semantically meaningful hole filling. Some use manually specified guidance, such as points of interest [16], lines [17], and perspective [18]. Other methods estimate image structures automatically, by various means such as tensor voting [19], search space constraints [20], statistics of patch offsets [21] and regularity in planar projection [22]. However, since they only depend on low-level visual features, such methods can only derive meaningful guidance in simple cases.

Hays and Efros [23] presented the first data-driven method to use a large reference dataset for hole filling. Whyte *et al.* [24] extended this approach with geometric and photometric registration. Zhu *et al.* [25] presented a faithful completion method for famous landmarks using their images found on the Internet. Barnes *et al.* [26] proposed a patch-based data structure for efficient patch query from the image database. These techniques may fail if the input image has a unique scene that cannot be found in the dataset.

Recently, deep learning emerges as a powerful tool for image completion. However, initial approaches can only handle very small regions [27]–[29]. Context encoders [30] introduced a generative adversarial loss [31] into the inpainting

network and combined this loss with an  $L_2$  pixel-wise reconstruction loss. It can produce plausible result in an  $128 \times 128$  image for a centered  $64 \times 64$  hole. Yang *et al.* [32] proposed to update the coarse result iteratively by searching nearest neural patches in the texture network for handling high resolution images. Yeh *et al.* [33] searched for the closest encoding in latent space with an adversarial loss and a weighted content loss, and then decoded it to a new complete image. More recently, Iizuka *et al.* [34] proposed an end-to-end completion network that was trained with a local and a global discriminator. The local discriminator examines the small region centered around the hole for local reality, while the global discriminator examines the entire image for global consistency. Dilated convolution layers [35] were also used to enlarge its spatial support. Yu *et al.* [36] extended this method with a novel contextual attention layer, which utilizes features surrounding the hole. Li *et al.* [37] focused on face completion with the help of semantic regularization. Existing learning-based methods are able to produce realistic completion results for general scenes such as landscapes and buildings, or certain specific objects such as human faces. However, they cannot handle portrait images as they do not consider the high-level semantic structures of human body.

### B. Human Parsing

Human parsing has been extensively studied in the literature. Liu *et al.* [38] combined CNN with KNN to predict matching confidence and displacements for regions into the test image. Co-CNN [1] integrates multi-level context into a unified network. Chen *et al.* [39] introduced an attention mechanism into the segmentation network, which learns to weight the importance of features at different scales and positions. Liang *et al.* [40], [41] employed Long Short-Term Memory (LSTM) units to exploit long range relationships in portrait images.

Without recovering the underlying human body structure, previous human parsing methods sometimes produce unreasonable results. In contrast, extracting human body structure has been extensively studied by pose estimation techniques [3], [4], [42]. Therefore, pose information has recently been integrated into human parsing frameworks to improve their performance. Gong *et al.* [2] presented a self-supervised structure-sensitive learning approach. It generates joint heatmaps from the parsing result map and the ground-truth label, and then calculates the joint structure loss by their Euclidean distances. JPPNet [43] builds a unified framework, which learns to predict human parsing maps and poses simultaneously. It then uses refinement networks to iteratively refine the parsing map and the pose. In our problem, pose estimation is helpful to guide the parsing prediction in the unknown region. Therefore, we adopt the basic architecture of JPPNet in our human parsing network, and specifically improve it for the following completion stage.

### C. Portrait Image Editing

Previous work has explored editing portrait images in various ways, and body shape editing is one that has been



Fig. 1. We address the problem of portrait image completion and extrapolation. (a) Shows that our method can remove the unwanted object from the portrait image. (b) Shows that our method can extrapolate the portrait image to recover the lower-body or the forehead.

paid special attention. Zhou *et al.* [44] integrated the 3D body morphable model into a single image warping approach, which is capable of reshaping a human body according to different weights or heights. PoseShop [45] constructed a large segmented human image database, from which new human figures can be synthesized with given poses. There are also methods for generating temporally-coherent human motion sequences with required poses [46] or shapes [47]. These techniques are mostly based on geometric deformation or database retrieval, so that their flexibility is limited.

Recently, many studies have been devoted to human image synthesis based on deep generative models. Lassner *et al.* [48] used the Variational Auto-Encoder (VAE) [49] to synthesize diverse clothes in human images. FashionGAN [50] can change the dress of the human figure in an input image with a given text description. Zhao *et al.* [51], Ma *et al.* [52] and Balakrishnan *et al.* [53] learned to generate human images with user-specified views or poses, while maintaining the character's identity and dress. Our goal is quite different from these methods, as we aim to recover only the missing parts of the human figure, not to generate an entirely new one.

### III. APPROACH

To realistically synthesize missing human body parts, one needs to estimate plausible region-level body structure as well as coherent textures in these regions. Training a network for simultaneously predicting both the structural configuration and appearance features is extremely difficult. We instead propose a deep learning framework which employs a two-stage approach to solve this problem. In stage-I, from the incomplete human body image we predict a complete parsing map through a human parsing network. In stage-II, we use a completion network to generate the inpainting result with the guidance of the parsing map, and a face refinement network to further improve the face region.

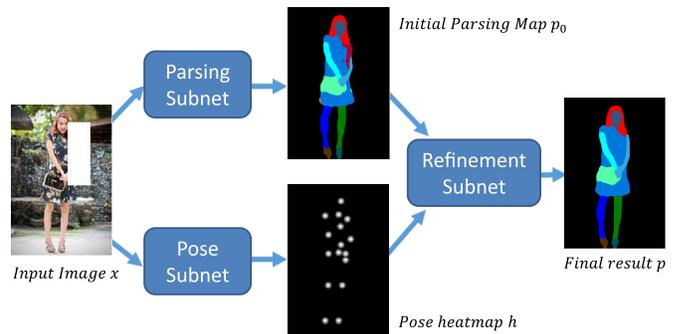


Fig. 2. Overview of our human parsing network at stage-I. The input image is fed into both the parsing subnet and the pose subnet. The two subnets then produce a human parsing map and pose heatmap respectively. Finally, the refinement network refines the parsing map with the help of the pose heatmap.

#### A. Stage-I: Human Parsing

Human parsing aims to segment a human body image into different parts with per-pixel labels. Compared with pose estimation, human parsing not only extracts the human body structure, but also estimates the image region boundary of each body part, which is beneficiary for image completion at the next stage. On the other hand, as suggested in JPPNet [43], pose estimation can help increase the accuracy of human parsing. Thus, we jointly train our human parsing network for both human body parsing and pose estimation. We then use the generated pose heatmap to refine the human parsing result within the unknown region.

The input of the human parsing network is an image with a pre-defined fill-in region. We denote the input image as  $x$ , the human parsing network as  $\mathbf{P}$ . Following JPPNet [43],  $\mathbf{P}$  consists of two subnets, a parsing subnet and a pose subnet. The two subnets share the first four stages in ResNet-101 [54]. In the parsing subnet, atrous spatial pyramid

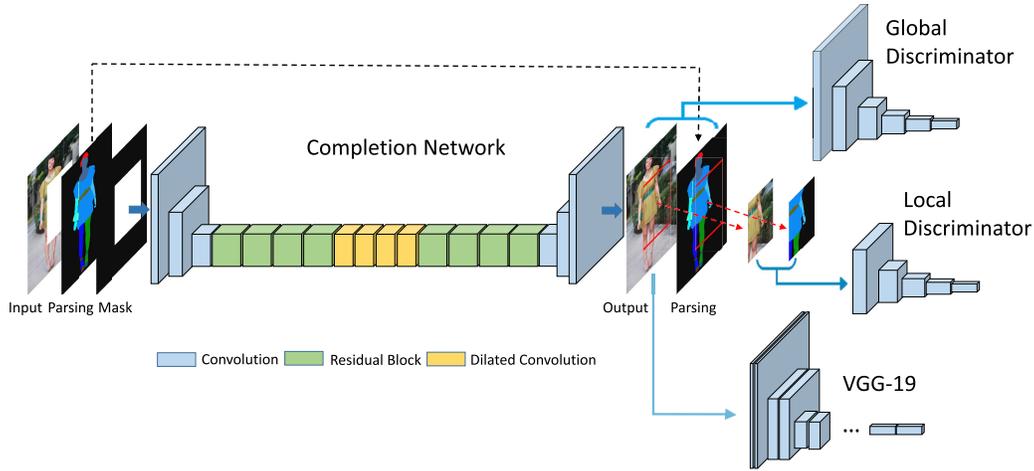


Fig. 3. Overview of our image completion networks at stage-II. The completion network generates the result image from the input image, parsing map and mask. It is trained with the perceptual loss and two adversarial losses. The perceptual loss measures feature maps from the VGG-19 network and the adversarial losses are backward from the discriminators. The result image with human parsing is fed into the discriminators, to tell whether it is realistic and in agreement with the structure. The input of global discriminator is the entire image with parsing, while the input of local discriminator is the local area surrounding the hole.

pooling (ASPP) [55] is applied to the fifth stage of ResNet-101 after the shared layers, to robustly segment different body parts. The parsing subnet produces the initial parsing result  $\mathbf{p}_0$ . In the pose subnet, several convolution layers are applied after the shared layers. The pose subnet produces the pose heatmap  $\mathbf{h}$ . Then a refinement subnet utilizes  $\mathbf{h}$  to refine the initial parsing result  $\mathbf{p}_0$  and produces the final result  $\mathbf{p}$ . We train the overall network end-to-end. Fig. 2 shows the overview of our human parsing network.

For effectiveness, we make several modifications to the network architecture of JPPNet [43] in  $\mathbf{P}$ . Specifically, we remove the pose refinement network in JPPNet and just keep the parsing refinement network. We thus do not apply iterative refinement and only refine the parsing map once. We regard  $\mathbf{p}$  as our final result instead of averaging all results over different iterations. After applying these simplifications, we have found that we can generate parsing results faster with better visual quality in the unknown region.

The typical loss function of human parsing is the mean of softmax losses at each pixel, formulated as:

$$L = \frac{1}{W \times H} \sum_i^W \sum_j^H L_S(p_{ij}, \hat{p}_{ij}). \quad (1)$$

$L_S$  denotes the softmax loss function.  $\hat{\mathbf{p}}$  denotes the ground-truth labels of the human parsing output.  $W$  and  $H$  are the width and height of the parsing map. Importantly, in our completion task we only need to generate image content inside the unknown hole region. Thus, the parsing accuracy inside the unknown region is much more important than that of the known region. Therefore, we propose a spatial weighted loss function that gives more weights on the pixels inside the hole. It is defined as:

$$L = \frac{1}{W \times H} \sum_i^W \sum_j^H (\alpha m_{ij} + 1) L_S(p_{ij}, \hat{p}_{ij}), \quad (2)$$

where  $\mathbf{m}$  is a binary mask indicating where to complete (1 for unknown pixels) and  $\alpha$  is a weighting parameter. We apply the spatial weighted loss to both the parsing subnet and the refinement subnet. We set  $\alpha = 9$  in our experiments by default.

### B. Stage-II: Image Completion

In stage-II, we use a completion network to synthesize missing regions with structure guidance. The input of the completion network consists of the input image  $\mathbf{x}$ , the human parsing map  $\mathbf{p}$  and the binary mask  $\mathbf{m}$  for marking the unknown region. We denote the completion network as  $\mathbf{G}$ , and its fully convolutional architecture is shown in Fig. 3. We bring in residual blocks [54] to enhance the representation ability, and employ dilated convolutions [35] to enlarge the spatial support. Instead of using a loss based on per-pixel color distance, we use a perceptual loss measured by feature maps from a pre-trained VGG-19 network [56]. Furthermore, to ensure that the network generates fine details that are semantically consistent with the known parts of the image, we feed the output of  $\mathbf{G}$  to local and global discriminators to measure adversarial losses. The details of the architecture and our training method are described in the following sections.

1) *Completion Network Architecture*: The completion network begins with a stride-1 convolutional layer and then uses two stride-2 convolutional layers to downsample the resolution to  $\frac{1}{4}$  of the input size. Four residual blocks are followed to extract features of both the input image and the human parsing map. Each residual block contains two stride-1 convolutional layers with a residual connection. Residual connection [54] has demonstrated its excellent ability to resolve the gradient vanishing problem, which usually happens in deep neural networks. In our experiments, we found residual blocks can significantly improve the quality of the synthesized results. Similar to Iizuka *et al.* [34], we use dilated convolutions [35] in the middle part of the network. Dilated convolution delivers large field of view without increasing computational cost,

which is necessary for global consistency of the completed image. After that, 4 residual blocks, 2 stride-2 deconvolutional layers and a stride-1 convolutional layer are applied sequentially. Kernel sizes of the first and the last convolution layers are 7. Kernel sizes of the other convolutional layers are all 3. Batch normalization layer and ReLU activation are applied after each convolution layer except the last one. We use a tanh layer in the end to normalize the result into  $[-1, 1]$ . The completion network  $\mathbf{G}$  is fully convolutional and can be adapted to arbitrary size of input images.

2) *Perceptual Loss*: Most previous completion methods [30], [34], [36] use  $L_1$  or  $L_2$  per-pixel distance as loss function to force the output image to be close to the input image. Recent works [57] have discovered that perceptual loss is more effective than  $L_1$  or  $L_2$  loss for synthesis tasks due to its advanced representation. Perceptual loss is first proposed by Gatys *et al.* [58] for image style transfer. Between a source and a target image, it measures the distance of their feature maps extracted from a pre-trained perception network. We use VGG-19 [56] as the perception network, which is pre-trained for ImageNet classification [59].

Denote the pre-trained perception network as  $\Phi$ . Layers in  $\Phi$  contain hierarchical features extracted from the input image. Shallower layers usually represent low-level features like colors and edges, while deeper layers represent high-level semantic features and more global statistics. Therefore, collection of layers from different levels contain rich information for image content synthesis. We define our perceptual loss function as:

$$L_p = \sum_{i=1}^n \|\Phi_i(\hat{x}) - \Phi_i(G(x, p, m))\|_2^2, \quad (3)$$

where  $n$  is the number of selected layers and  $\Phi_i$  is the  $i$ -th selected layer. In our experiments,  $\Phi$  includes *relu1\_2*, *relu2\_2*, *relu3\_2*, *relu4\_2* and *relu5\_2* layers in VGG-19 [56].

3) *Discriminators*: Using only perceptual loss in training often leads to obvious artifacts in output images. Inspired by previous image completion methods [30], [34], [36], we add an adversarial loss into the completion network to prevent generating unrealistic results. The adversarial loss is based on Generative Adversarial Networks (GANs) [31]. GANs consist of two networks, a generator and a discriminator. The discriminator learns to distinguish real images from generated ones. The generator tries to produce realistic images to fool the discriminator. The two networks compete with each other, and after convergence the generator can produce realistic output.

In our task, the generated result (i.e. the filled holes) needs to be not only realistic, but also coherent with the structure guidance. Therefore, we employ conditional GANs [60] in our completion network, where the discriminator  $\mathbf{D}$  also learns to determine whether the generated content conforms to the condition, i.e. the human parsing map  $\mathbf{p}$ . Hence we define the adversarial loss as:

$$L_{adv} = \min_G \max_D \mathbb{E}[\log D(p, \hat{x}) + \log(1 - D(p, G(x, p, m)))]. \quad (4)$$

Inspired by Iizuka *et al.* [34], we use both global and local discriminators in our completion network. The input to the global discriminator is a concatenation of the input image and the parsing map. The input to the local discriminator is the image region and its corresponding parsing patch centered around the hole to be completed. At training stage, the size of the image is scaled to  $256 \times 256$  and the size of the patch is fixed to  $128 \times 128$ . As in Iizuka *et al.* [34], we only use  $\mathbf{G}$  to complete a single hole at training time. Note that  $\mathbf{G}$  can fill multiple holes at once in testing stage. Noted by Pinheiro *et al.* [61], unequal channels of the signal may cause imbalance that one signal dominates the other ones. Therefore, we repeat the one-channel parsing map for three times to match the RGB channels. We adopt the network architecture proposed by DCGAN [62] for two discriminators. All convolutional layers in the discriminators have kernels of size 4 and stride 2, except the last stride-1 convolution. The global discriminator has 7 convolution layers and the local discriminator has 6. They both end with a Sigmoid layer to produce a true or false label.

We define the overall loss function for the completion network as

$$L_c = \lambda_p L_p + \lambda_g L_{adv-g} + \lambda_l L_{adv-l}. \quad (5)$$

$L_{adv-g}$  and  $L_{adv-l}$  represent adversarial losses for the global and the local discriminators, respectively.  $\lambda_p$ ,  $\lambda_g$  and  $\lambda_l$  are the hyper parameters to balance the different losses. We set  $\lambda_p = 100$ ,  $\lambda_g = \lambda_l = 1$  in our experiments.

### C. Face Refinement

The human face may only occupy a small area in the input image, but contains many delicate details that the human vision system is sensitive to. Completion network  $\mathbf{G}$  can recover general missing regions well, but may have difficulties with human faces as it is not specifically trained on them. Similar to Chan *et al.* [63], we thus propose a dedicated face refinement network  $\mathbf{F}$  to refine inpainted human faces.

We crop the inpainting result  $G(x, p, m)$  produced by the completion network to a window around the synthesized face  $G_f(x, p, m)$ , then feed it along with the cropped parsing map  $\mathbf{p}_f$  and the cropped mask  $\mathbf{m}_f$  into the face network  $\mathbf{F}$ . We calculate the centre of mass of the face region through the human parsing map and crop the region around the center by the size of  $64 \times 64$ . Then  $\mathbf{F}$  produces the residual face image  $R_f = F(G_f(x, p, m), \mathbf{p}_f, \mathbf{m}_f)$  with the size of  $64 \times 64$ . The final synthesized result  $\hat{\mathbf{R}}$  is the addition of the residual face image  $\mathbf{R}_f$  and the initial inpainting result  $G(x, p, m)$ . Fig. 4 shows the pipeline of the face refinement.

After the training of the completion network is finished, we train the face network with the parameters of the completion network fixed. For the realism of the refined face image, we introduce a face discriminator  $\mathbf{D}_f$  to distinguish generated faces from the real ones. We also use the perceptual loss to ensure that the generated faces are perceptually indistinguishable from the ground-truth. The architecture of the face network  $\mathbf{F}$  is similar to the completion network  $\mathbf{G}$  except

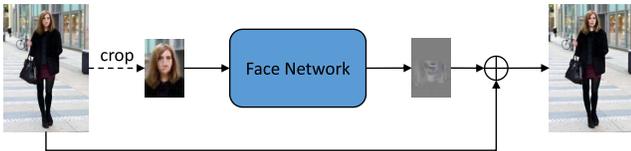


Fig. 4. Overview of our face refinement network. We crop the initial inpainting result of the face region and feed it into the face network. The final output is the sum of the produced residual image and the initial result.

that the number of the residual blocks is reduced to 4. The face discriminator  $\mathbf{D}_f$  also adopts the architecture proposed by DCGAN [62]. The full objective is

$$L_f = \lambda \sum_{i=1}^n \left\| \Phi_i(\hat{x}_f) - \Phi_i(\hat{R}_f) \right\|_2^2 + \log(1 - D_f(p_f, \hat{R}_f)). \quad (6)$$

We set  $\lambda = 10$  in our experiments.

#### D. Implementation Details

We train the networks at two stages separately. At stage-I, we train the human parsing network by stochastic gradient descent with momentum. We set the learning rate to 0.0001 and momentum to 0.9. At stage-II, We use the Adam solver [64] with a batch size of 1 to train all the networks. We set the learning rate to 0.0002. We scale the input image to  $256 \times 256$  at training time. We randomly crop a rectangular region in the input image. The edge of the rectangle is randomly set in range [64, 128]. Since we are concerned about completion for the human body, we make sure that the rectangle overlaps the human body in image. We set the pixels inside the cropped region to the mean pixel value of the datasets. To prevent overfitting, we apply several data augmentation methods for training, including scaling, cropping and left-right flipping. When training the face network, we randomly crop a rectangular region in the face area by the size of  $32 \times 32$ .

At testing time, the input image goes through the human parsing network, the completion network and the face network sequentially to generate the final result. Because all the networks are fully-convolutional, our method can be adapted to arbitrary size of image. We use Poisson image blending [65] to post-process the final result, as previous completion methods [32], [34] do, for more consistent boundary of the completion region.

## IV. RESULTS

We evaluate our method on two human image datasets, ATR dataset [1] and LIP dataset [2]. ATR dataset contains 17700 human images with parsing annotations on 17 body part categories. We randomly select 1/10 of the images for testing and others for training. LIP dataset contains 30462 training images and 10000 validation images. Each image in LIP dataset has per-pixel parsing annotations on 19 categories and annotations on 16 keypoints for the pose estimation. Because there is no pose annotation in ATR dataset, we use JPPNet [43] pre-trained on LIP dataset to predict the body pose for each image in ATR dataset and use the result as the ground-truth label for training.

TABLE I  
COMPARISON WITH EXISTING COMPLETION METHODS  
IN TERMS OF  $L_1$  ERROR, PSNR, SSIM AND FID

Dataset	Method	$L_1$ error	PSNR	SSIM	FID
ATR	PatchMatch [13]	43.0182	18.8905	0.8615	50.9721
	Iizuka et al. [34]	12.5181	23.7539	0.9532	21.4838
	Yu et al. [36]	13.7748	23.6954	0.9481	16.4192
	G	10.4527	25.2684	0.9591	11.8323
	G+P	9.3675	27.0852	<b>0.9736</b>	<b>10.6247</b>
	G+P+F(ours)	<b>8.4620</b>	<b>27.1511</b>	0.9730	10.6578
LIP	PatchMatch [13]	35.4193	20.6823	0.9063	10.9311
	Iizuka et al. [34]	12.8931	25.0089	0.9646	9.3290
	Yu et al. [36]	16.8708	22.9498	0.9437	5.3085
	G	12.9677	25.1877	0.9630	3.8510
	G+P	10.6944	26.0007	<b>0.9696</b>	3.1399
	G+P+F(ours)	<b>10.3144</b>	<b>26.0641</b>	0.9693	<b>2.9604</b>

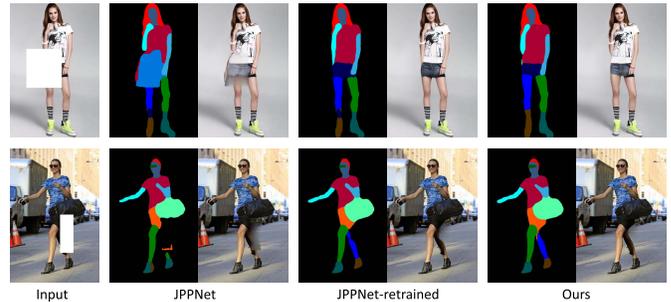


Fig. 5. Completion results by different human parsing methods. We complete the input image with different human parsing method. We show that our method leads to the best completion performance.

We initialize the human parsing network with the parameters loaded from pre-trained JPPNet and we then train the network for 10 epochs on the two datasets. Meanwhile, we train the completion network for 200 epochs on ATR dataset and 100 epochs on LIP dataset. After that, we train the face network with 100 epochs on ATR dataset and 50 epochs on LIP dataset. It takes about 20 hours to train the human parsing network, 100 hours for the completion network and 50 hours for the face network. When testing on GPU, our two-stage framework spends about 0.48s to produce the final result for an input image of size  $256 \times 256$ . We evaluate our method on using an Intel(R) Core(TM) i7-4770 CPU @ 3.40GHz with 4 cores and NVidia GTX 1080 Ti GPU.

#### A. Comparisons With Existing Completion Methods

We compare our approach with several existing completion methods, including Photoshop Content Aware Fill (PatchMatch) [13], Iizuka *et al.* [34] and Yu *et al.* [36]. We train the model of Iizuka *et al.* [34] on training sets of ATR and LIP for the same epoch numbers as our method. We initialize the model with the parameters pre-trained on Places2 dataset [66]. We also train the model of Yu *et al.* [36] for the same epoch and initialize it with the model pre-trained on ImageNet dataset. We test the models of Iizuka *et al.* [34], Yu *et al.* [36] and our method at the scale of  $256 \times 256$ . We then resize the result to the scale of the original image for a fair comparison. Fig. 7 shows some comparisons from the test set in ATR.

TABLE II

COMPARISON OF THE HUMAN PARSING METHODS IN TERMS OF MEAN IOU. WE MEASURE THE IOU FOR BOTH THE ENTIRE IMAGES AND FOR THE UNKNOWN REGIONS ON ATR DATASET AND LIP DATASET

Dataset	Method	Entire image	Unknown region
ATR	JPPNet [43]	0.5193	0.1132
	JPPNet-retrained	0.5786	0.3656
	Ours ( $\alpha = 0$ )	<b>0.5965</b>	0.3735
	Ours	0.5937	<b>0.3971</b>
LIP	JPPNet [43]	0.4125	0.0796
	JPPNet-retrained	0.4649	0.3099
	Ours ( $\alpha = 0$ )	<b>0.4817</b>	0.3257
	Ours	0.4685	<b>0.3381</b>

Fig. 8 shows some comparison results from the validation set in LIP. These results suggest that our approach can generate plausible results for large holes, while other methods fail to handle complicated human body structures. Fig. 9 also shows some completion results with multiple holes by our method.

We also compare our approach with other approaches quantitatively. We report the evaluation in terms of  $L_1$  error, PSNR and SSIM [67]. Table. I shows that our method outperforms other methods under these measurements on the two public datasets. Please note that these evaluation metrics are by no means the best metrics for evaluating the results, since there may be many visually pleasing completion results that are equally acceptable other than the ground truth images. Recently, FID [68] was proposed to measure the perceptual quality of the generative model. Our method also achieves the best performance under this perceptual metric.

### B. Comparisons on Human Parsing

We compare our human parsing network with JPPNet [43] to show its effectiveness for completion. We use JPPNet trained on complete images and refined on incomplete images with the same pattern as our parsing network. We apply the mean of intersection over union (IoU) as the performance metric. We measure the mean IoU for the entire image and the unknown region. We report the results in Table. II. As shown in the table, our method achieves the best parsing performance especially inside the unknown region. Fig. 5 also shows that our human parsing method leads to the best completion result.

### C. Completion Results for Animals

Although we focus on human body completion in this paper, the proposed learning framework is generic and can be applied to other types of images with highly structured visual features, such as animal images. To demonstrate it, we train our model on the Horse-Cow dataset [69]. Horse-Cow dataset contains 295 images for training and 277 images for testing. We additionally label the keypoints of poses for this dataset since it only includes parsing annotations. We re-train the parsing network for 300 epochs and the completion network for 1000 epochs on the training set. Notably, we train the horses and cows together since these two kinds of animals have similar structures. Fig. 10 shows some completion results using the testing images. We also compare our method with



Fig. 6. Completion results before (middle) and after (right) the face refinement. The face network refines the facial structures and makes the completion result more realistic.

Iizuka *et al.* [34] and Yu *et al.* [36]. We train the models of the two methods for the same epochs as ours. From the results, we can see that our two-stage deep structure can be generalized to other images in which the inherent semantic structures could be captured and predicted.

### D. Ablation Study

We perform an ablation study to validate the importance of our human body structure guidance and face refinement network. We train an exactly the same completion network without using the parsing map as input. The training protocols are kept unchanged as well. This model is denoted by G. We also train a completion network with human body parsing maps, denoted as G+P. Similarly, our full model with the face refinement network is denoted as G+P+F. Results in Table. I show that the parsing guidance and the face refinement are both beneficiary to the completion result. Fig. 6 compares the completion results before and after the face refinement, showing the effectiveness of our face network.

We also evaluate the effectiveness of our spatial weighted loss for human parsing, by comparing to the results when  $\alpha$  is 0. As shown in Table. II, the spatial weighted loss can increase the parsing accuracy inside the holes, which is good for the following completion process.

## V. APPLICATIONS

### A. Portrait Extrapolation

As a general photography rule, it is not recommended to cut off people's feet or foreheads in the composition. Nevertheless, amateur photographers often make such mistakes. It is desirable to extend the portrait image to recover the missing feet or faces to create a better composition. Fortunately, this problem can be directly solved by our framework, by treating it as a completion problem, where the unknown region is at the bottom or top part of the extrapolated image.

We re-train the entire framework by fixing the cropped rectangle in the input image. For the downward extrapolation, the bottom is cropped by the size of  $64 \times 256$ , and for the upward extrapolation, the top is cropped by the size of  $32 \times 256$ . The reason for the different heights is that normally the missing head region is shorter than the missing

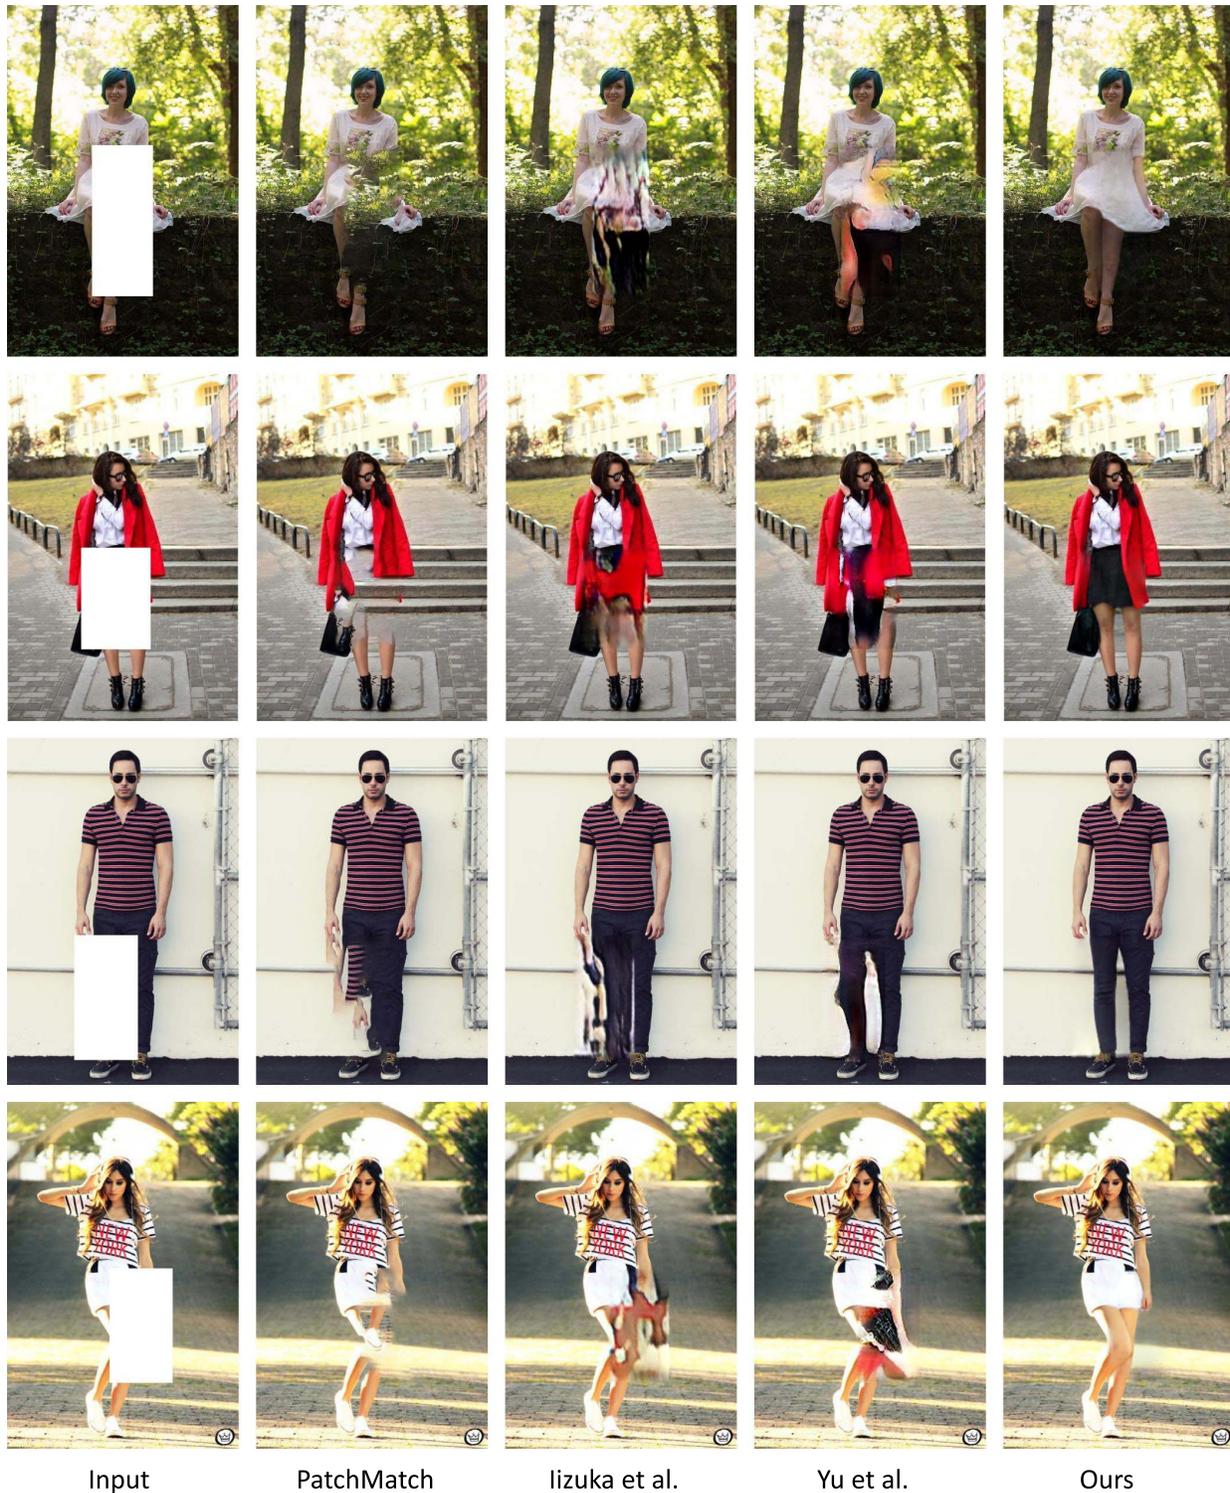


Fig. 7. Comparison results on the ATR testing set. We compare our method with PatchMatch [13], Iizuka *et al.* [34] and Yu *et al.* [36].

legs/feet region. We abandon the local discriminator due to the exceeded scale of the hole. We also add dropout layers into the completion network for more flexible results. Our human parsing network can analyze the posture and predict reasonable structures for the legs and shoes, as well as the hair and face. Then the completion network can generate the

extrapolated image that contains consistent body parts and background. The face region would be further enhanced by the final face network. For comparisons, we also train the model of Yu *et al.* [36] with the same pattern. Note that, we also remove the local discriminator in their framework. To demonstrate the importance of the structural information in

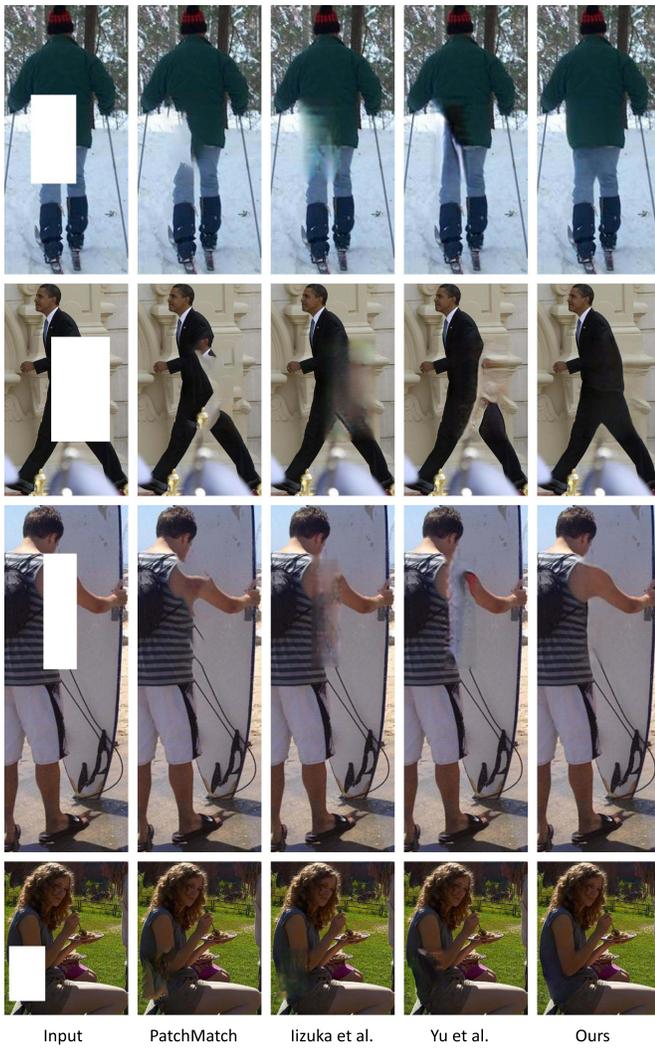


Fig. 8. Comparison results on the LIP validation set. We compare our method with PatchMatch [13], lizuka *et al.* [34] and Yu *et al.* [36].



Fig. 9. Completion results with multiple holes on the ATR testing set.

the portrait extrapolation task, we train the model G without the parsing guidance as a comparison, as shown in Fig. 11.

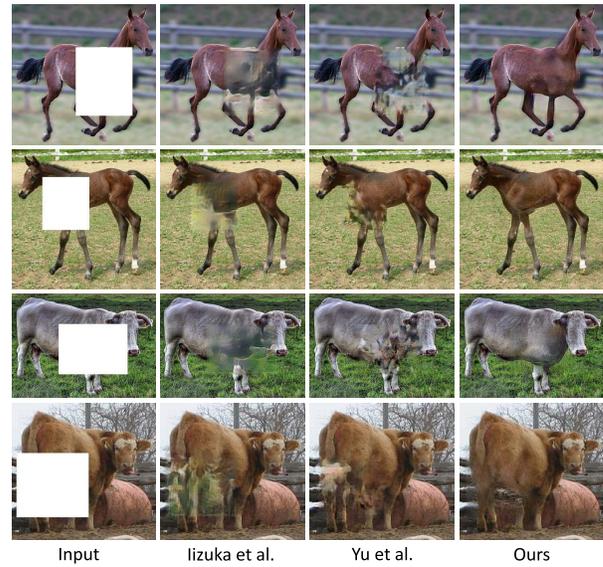


Fig. 10. Comparison results on Horse-Cow dataset. We compare our method with lizuka *et al.* [34] and Yu *et al.* [36]. We show that our method is generic and also available for animals like horse or cow.

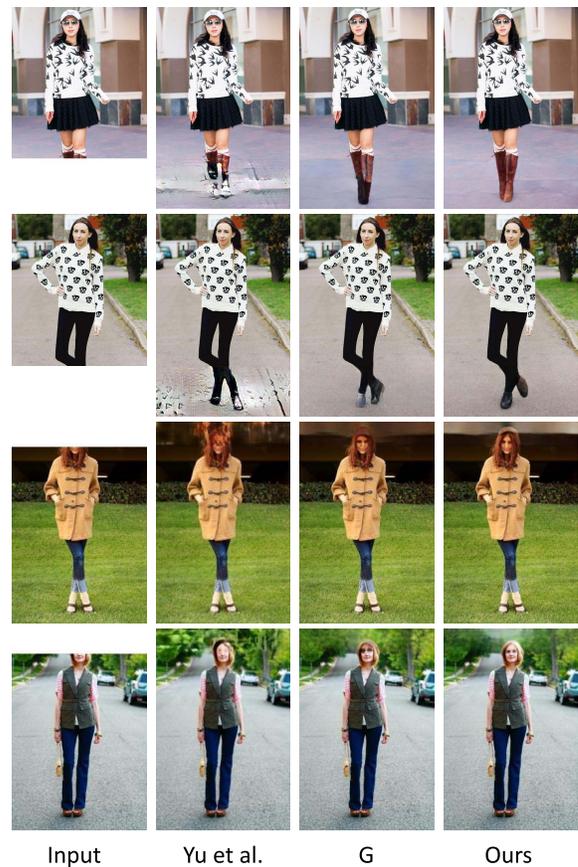


Fig. 11. Comparison results of portrait extrapolation. We compare our full model with Yu *et al.* [36] and the model G. The results show that our method can generate plausible lower body parts and faces, while other methods fail.

These results suggest that by generating high fidelity human body completion results, our method opens new possibilities for portrait image enhancement and editing.

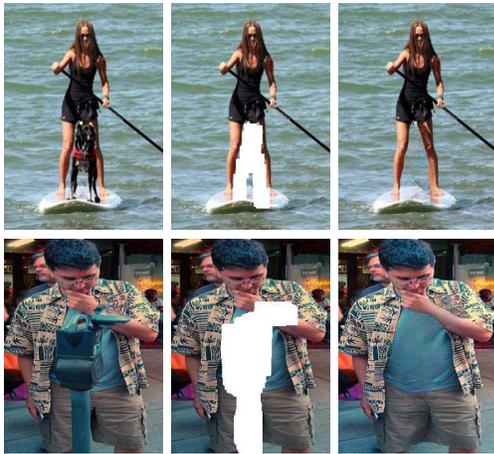


Fig. 12. Examples of occlusion removal by our approach.



Fig. 13. Some failure cases from the ATR dataset. Our method cannot produce satisfactory result when the textured region is heavily covered and will be confused with multiple persons.

### B. Occlusion Removal

Occlusion removal is a natural application for the task of image completion. Fig. 12 shows that our approach can recover the full human body when removing the unwanted objects. Because our framework has a human parsing stage to generate correct structural information for the hole region, we can handle large occlusions as shown in the bottom example in Fig. 12. Note that we preserve the geometric features along the arm and clothes boundaries after completion.

## VI. CONCLUSION AND LIMITATION

We propose a two-stage deep learning framework to solve portrait image completion problem. We first employ a human parsing network to extract structural information from the input image. Then we employ a completion network to generate the unknown region with the guidance of the parsing result and a following face network to refine the face appearance. We have demonstrated that, aware of the structure of the human body, we can produce more reasonable and more realistic result compared to other methods. And we have shown the capability of our method for applications like occlusion removal and portrait extrapolation. Besides humans, we have also experimented our method on animals and achieved impressive completion results, which indicates that our framework can be extended to other kinds of images where the inherent semantic structures could be encoded and predicted.

Our method may fail in some cases as shown in Fig. 13. Firstly, when most of the textured region or the logo is covered, our method may not yield satisfactory completion result due to the inadequate information, as shown in Fig. 13 (left). Secondly, our model will be confused when the hole region is connected to multiple persons, as shown in Fig. 13 (right). That is because our current human parsing network is only trained by portrait images with a single person. We plan to extend our framework to deal with multiple persons in the future.

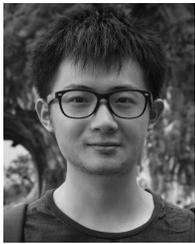
### ACKNOWLEDGMENT

The authors would like to thank all the reviewers.

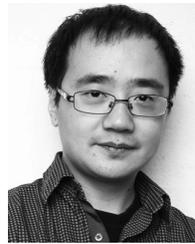
### REFERENCES

- [1] X. Liang *et al.*, “Human parsing with contextualized convolutional neural network,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1386–1394.
- [2] K. Gong, X. Liang, D. Zhang, X. Shen, and L. Lin, “Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 932–940.
- [3] A. Newell, K. Yang, and J. Deng, “Stacked hourglass networks for human pose estimation,” in *Proc. Eur. Conf. Comput. Vis.* Springer, 2016, pp. 483–499. [Online]. Available: [https://rd.springer.com/chapter/10.1007/978-3-319-46484-8\\_29](https://rd.springer.com/chapter/10.1007/978-3-319-46484-8_29)
- [4] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, “Realtime multi-person 2D pose estimation using part affinity fields,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 7291–7299.
- [5] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, “Image inpainting,” in *Proc. 27th Annu. Conf. Comput. Graph. Interact. Techn.*, 2000, pp. 417–424.
- [6] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, and J. Verdera, “Filling-in by joint interpolation of vector fields and gray levels,” *IEEE Trans. Image Process.*, vol. 10, no. 8, pp. 1200–1211, Aug. 2001.
- [7] M. Bertalmio, L. Vese, G. Sapiro, and S. Osher, “Simultaneous structure and texture image inpainting,” *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 882–889, Aug. 2003.
- [8] A. A. Efros and T. K. Leung, “Texture synthesis by non-parametric sampling,” in *Proc. 7th IEEE Int. Conf. Comput. Vis.*, Sep. 1999, pp. 1033–1038.
- [9] A. A. Efros and W. T. Freeman, “Image quilting for texture synthesis and transfer,” in *Proc. 28th Annu. Conf. Comput. Graph. Interact. Techn.*, 2001, pp. 341–346.
- [10] V. Kwatra, A. SchÖdl, I. Essa, G. Turk, and A. Bobick, “Graphcut textures: Image and video synthesis using graph cuts,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 277–286, 2003.
- [11] A. Criminisi, P. Pérez, and K. Toyama, “Region filling and object removal by exemplar-based image inpainting,” *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1200–1212, Sep. 2004.
- [12] D. Simakov, Y. Caspi, E. Shechtman, and M. Irani, “Summarizing visual data using bidirectional similarity,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.
- [13] C. Barnes, E. Shechtman, A. Finkelstein, and D. Goldman, “PatchMatch: A randomized correspondence algorithm for structural image editing,” *ACM Trans. Graph.*, vol. 28, no. 3, p. 24, 2009.
- [14] C. Barnes, E. Shechtman, D. B. Goldman, and A. Finkelstein, “The generalized Patchmatch correspondence algorithm,” in *Proc. Eur. Conf. Comput. Vis.* Springer, 2010, pp. 29–43. [Online]. Available: [https://rd.springer.com/chapter/10.1007/978-3-642-15558-1\\_3](https://rd.springer.com/chapter/10.1007/978-3-642-15558-1_3)
- [15] S. Darabi, E. Shechtman, C. Barnes, D. B. Goldman, and P. Sen, “Image melding: Combining inconsistent images using patch-based synthesis,” *ACM Trans. Graph.*, vol. 31, no. 4, pp. 1–82, 2012.
- [16] I. Drori, D. Cohen-Or, and H. Yeshurun, “Fragment-based image completion,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 303–312, 2003.
- [17] J. Sun, L. Yuan, J. Jia, and H.-Y. Shum, “Image completion with structure propagation,” *ACM Trans. Graph.*, vol. 24, no. 3, pp. 861–868, 2005.
- [18] D. Pavić, V. Schönefeld, and L. Kobbelt, “Interactive image completion with perspective correction,” *Vis. Comput.*, vol. 22, nos. 9–11, pp. 671–681, 2006.

- [19] J. Jia and C.-K. Tang, "Image repairing: Robust image synthesis by adaptive ND tensor voting," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, vol. 1, Jun. 2003, pp. 1–1.
- [20] J. Kopf, W. Kienzle, S. Drucker, and S. B. Kang, "Quality prediction for image completion," *ACM Trans. Graph.*, vol. 31, no. 6, p. 131, 2012.
- [21] K. He and J. Sun, "Statistics of patch offsets for image completion," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2012, pp. 16–29. [Online]. Available: [https://rd.springer.com/chapter/10.1007/978-3-642-33709-3\\_2](https://rd.springer.com/chapter/10.1007/978-3-642-33709-3_2)
- [22] J.-B. Huang, S. B. Kang, N. Ahuja, and J. Kopf, "Image completion using planar structure guidance," *ACM Trans. Graph.*, vol. 33, no. 4, p. 129, 2014.
- [23] J. Hays and A. A. Efros, "Scene completion using millions of photographs," *ACM Trans. Graph.*, vol. 26, no. 3, p. 4, 2007.
- [24] O. Whyte, J. Sivic, and A. Zisserman, "Get out of my picture! Internet-based inpainting," in *Proc. BMVC*, 2009, vol. 2, no. 4, p. 5.
- [25] Z. Zhu, H.-Z. Huang, Z.-P. Tan, K. Xu, and S.-M. Hu, "Faithful completion of images of scenic landmarks using Internet images," *IEEE Trans. Vis. Comput. Graphics*, vol. 22, no. 8, pp. 1945–1958, Aug. 2016.
- [26] C. Barnes, F.-L. Zhang, L. Lou, X. Wu, and S.-M. Hu, "PatchTable: Efficient patch queries for large datasets and applications," *ACM Trans. Graph.*, vol. 34, no. 4, p. 97, 2015.
- [27] J. Xie, L. Xu, and E. Chen, "Image denoising and Inpainting with deep neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 341–349.
- [28] R. Köhler, C. Schuler, B. Schölkopf, and S. Harmeling, "Mask-specific Inpainting with deep neural networks," in *Proc. German Conf. Pattern Recognit.* Springer, 2014, pp. 523–534.
- [29] J. S. Ren, L. Xu, Q. Yan, and W. Sun, "Shepard convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 901–909.
- [30] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context encoders: Feature learning by inpainting," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2536–2544.
- [31] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [32] C. Yang, X. Lu, Z. Lin, E. Shechtman, O. Wang, and H. Li, "High-resolution image inpainting using multi-scale neural patch synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, vol. 1, no. 2, pp. 6721–6729.
- [33] 33 R. A. Yeh, C. Chen, T. Y. Lim, A. G. Schwing, M. Hasegawa-Johnson, and M. N. Do, "Semantic image inpainting with deep generative models," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 5485–5493.
- [34] S. Iizuka, E. Simo-Serra, and H. Ishikawa, "Globally and locally consistent image completion," *ACM Trans. Graph.*, vol. 36, no. 4, p. 107, 2017.
- [35] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," 2015, *arXiv:1511.07122*. [Online]. Available: <https://arxiv.org/abs/1511.07122>
- [36] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Generative image inpainting with contextual attention," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2018, pp. 5505–5514.
- [37] Y. Li, S. Liu, J. Yang, and M.-H. Yang, "Generative face completion," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3911–3919.
- [38] S. Liu *et al.*, "Matching-CNN meets KNN: Quasi-parametric human parsing," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1419–1427.
- [39] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille, "Attention to scale: Scale-aware semantic image segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3640–3649.
- [40] X. Liang, X. Shen, D. Xiang, J. Feng, L. Lin, and S. Yan, "Semantic object parsing with local-global long short-term memory," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3185–3193.
- [41] X. Liang, X. Shen, J. Feng, L. Lin, and S. Yan, "Semantic object parsing with graph LSTM," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2016, pp. 125–143. [Online]. Available: [https://rd.springer.com/chapter/10.1007/978-3-319-46448-0\\_8](https://rd.springer.com/chapter/10.1007/978-3-319-46448-0_8)
- [42] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh, "Convolutional pose machines," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 4724–4732.
- [43] X. Liang, K. Gong, X. Shen, and L. Lin, "Look into person: Joint body parsing & pose estimation network and a new benchmark," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 4, pp. 871–885, Apr. 2018.
- [44] S. Zhou, H. Fu, L. Liu, D. Cohen-Or, and X. Han, "Parametric reshaping of human bodies in images," *ACM Trans. Graph.*, vol. 29, no. 4, p. 126, 2010.
- [45] T. Chen, P. Tan, L.-Q. Ma, M.-M. Cheng, A. Shamir, and S.-M. Hu, "PoseShop: Human image database construction and personalized content synthesis," *IEEE Trans. Vis. Comput. Graphics*, vol. 19, no. 5, pp. 824–837, May 2013.
- [46] F. Xu *et al.*, "Video-based characters: Creating new human performances from a multi-view video database," *ACM Trans. Graph.*, vol. 30, no. 4, p. 32, Jul. 2011.
- [47] A. Jain, T. Thormählen, H.-P. Seidel, and C. Theobalt, "MovieReshape: Tracking and reshaping of humans in videos," *ACM Trans. Graph.*, vol. 29, no. 6, p. 148, 2010.
- [48] C. Lassner, G. Pons-Moll, and P. V. Gehler, "A generative model of people in clothing," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 853–862.
- [49] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*. [Online]. Available: <https://arxiv.org/abs/1312.6114>
- [50] S. Zhu, R. Urtasun, S. Fidler, D. Lin, and C. C. Loy, "Be your own Prada: Fashion synthesis with structural coherence," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 1680–1688.
- [51] B. Zhao, X. Wu, Z.-Q. Cheng, H. Liu, and J. Feng, "Multi-view image generation from a single-view," 2017, *arXiv:1704.04886*. [Online]. Available: <https://arxiv.org/abs/1704.04886>
- [52] L. Ma, X. Jia, Q. Sun, B. Schiele, T. Tuytelaars, and L. Van Gool, "Pose guided person image generation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 405–415.
- [53] G. Balakrishnan, A. Zhao, A. V. Dalca, F. Durand, and J. Guttag, "Synthesizing images of humans in unseen poses," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2018, pp. 8340–8348.
- [54] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [55] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, Apr. 2017.
- [56] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*. [Online]. Available: <https://arxiv.org/abs/1409.1556>
- [57] Q. Chen and V. Koltun, "Photographic image synthesis with cascaded refinement networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, vol. 1, Oct. 2017, pp. 1511–1520.
- [58] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2414–2423.
- [59] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [60] M. Mirza and S. Osindero, "Conditional generative adversarial nets," 2014, *arXiv:1411.1784*. [Online]. Available: <https://arxiv.org/abs/1411.1784>
- [61] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollár, "Learning to refine object segments," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2016, pp. 75–91. [Online]. Available: [https://rd.springer.com/chapter/10.1007/978-3-319-46448-0\\_5](https://rd.springer.com/chapter/10.1007/978-3-319-46448-0_5)
- [62] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," 2015, *arXiv:1511.06434*. [Online]. Available: <https://arxiv.org/abs/1511.06434>
- [63] C. Chan, S. Ginosar, T. Zhou, and A. A. Efros, "Everybody dance now," 2018, *arXiv:1808.07371*. [Online]. Available: <https://arxiv.org/abs/1808.07371>
- [64] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <https://arxiv.org/abs/1412.6980>
- [65] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 313–318, 2003.
- [66] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba, "Places: A 10 million image database for scene recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 6, pp. 1452–1464, Jun. 2017.
- [67] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [68] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local Nash equilibrium," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 6626–6637.
- [69] J. Wang and A. L. Yuille, "Semantic part segmentation using compositional model combining shape and appearance," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1788–1797.



**Xian Wu** received the B.S. degree from Tsinghua University in 2015, where he is currently pursuing the Ph.D. degree. His research interests include image synthesis/editing and deep learning in computer graphics.



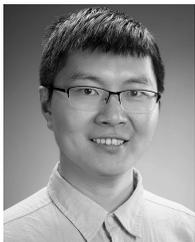
**Jue Wang** (SM'12) received the B.E. and M.Sc. degrees from the Department of Automation, Tsinghua University in 2000 and 2003, respectively, and the Ph.D. degree in electrical engineering from the University of Washington in 2007. He has been a Principle Research Scientist with Adobe Research for nine years. He is currently the Director of Megvii (Face++) Research USA. His research interests include image and video processing, computational photography, computer graphics, and vision. He is a member of ACM. He received the Microsoft Research Fellowship and the Yang Research Award from the University of Washington in 2006.



**Rui-Long Li** received the B.S. degree from Tsinghua University in 2016, where he is currently pursuing the master's degree. His research interests include computer vision and image/video processing.



**Ariel Shamir** received the B.Sc. and M.Sc. degrees (*cum laude*) in math and computer science from Hebrew University, Jerusalem, and the Ph.D. degree in computer science in 2000. He spent two years as a Postdoctoral Fellow with the Computational Visualization Center, The University of Texas at Austin, Austin, TX, USA. He was a Visiting Scientist with the Mitsubishi Electric Research Labs, Cambridge, MA, USA, in 2006, and Disney Research in 2014. He is currently the Dean of the Efi Arazi School of Computer Science, Interdisciplinary Center Herzliya, Israel. His current research interests include geometric modeling, computer graphics, fabrication, visualization, and machine learning. He is a member of the ACM SIGGRAPH, and the IEEE Computer and Eurographics societies.



**Fang-Lue Zhang** (M'13) received the bachelor's degree from Zhejiang University in 2009 and the Ph.D. degree from Tsinghua University in 2015. He is currently a Lecturer with the Victoria University of Wellington. His research interests include image and video editing, computer vision, and computer graphics. He is a member of ACM.



**Shi-Min Hu** received the Ph.D. degree from Zhejiang University in 1996. He is currently a Professor with the Department of Computer Science and Technology, Tsinghua University, Beijing. He has published over 100 articles in journals and refereed conference. His research interests include digital geometry processing, video processing, rendering, computer animation, and computer-aided geometric design. He is on the editorial board of several journals, including the IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS, *Computer Aided Design*, and *Computer and Graphics*. He is the Editor-in-Chief of *Computational Visual Media*.



**Jian-Cheng Liu** is currently pursuing the bachelor's degree with Tsinghua University. His research interests include computer vision and machine learning.