

Recovering Faces from Portraits with Auxiliary Facial Attributes

Fatemeh Shiri¹, Xin Yu¹, Fatih Porikli¹, Richard Hartley^{1,2}, Piotr Koniusz^{2,1}

¹Australian National University, ²Data61/CSIRO

firstname.lastname@{anu.edu.au¹, data61.csiro.au²}

Abstract

Recovering a photorealistic face from an artistic portrait is a challenging task since crucial facial details are often distorted or completely lost in artistic compositions. To handle this loss, we propose an Attribute-guided Face Recovery from Portraits (AFRP) that utilizes a Face Recovery Network (FRN) and a Discriminative Network (DN). FRN consists of an autoencoder with residual block-embedded skip-connections and incorporates facial attribute vectors into the feature maps of input portraits at the bottleneck of the autoencoder. DN has multiple convolutional and fully-connected layers, and its role is to enforce FRN to generate authentic face images with corresponding facial attributes dictated by the input attribute vectors. For the preservation of identities, we impose the recovered and ground-truth faces to share similar visual features. Specifically, DN determines whether the recovered image looks like a real face and checks if the facial attributes extracted from the recovered image are consistent with given attributes. Our method can recover photorealistic identity-preserving faces with desired attributes from unseen stylized portraits, artistic paintings, and hand-drawn sketches. On large-scale synthesized and sketch datasets, we demonstrate that our face recovery method achieves state-of-the-art results.

1. Introduction

Numerous style transfer methods have been proposed to transfer arbitrary artwork styles into content images. In contrast to image stylization, we address a challenging inverse problem called photorealistic face recovery from stylized portraits which aims at recovering a photorealistic face image from a given stylized portrait. The recovery of the latent photorealistic face from its artistic portrait can help facial analysis and the digital entertainment. Facial details in stylized portraits contain artistic effects and distortions such as profile edges and texture changes as shown in Fig. 1(b). These artistic effects result in a partial loss of facial details and identity-related information. Moreover, stylized face images may contain various facial expressions, facial distortions and misalignments. Off-the-shelf facial landmark detectors often fail to localize facial landmarks correctly as

shown in Fig. 1(c). Thus, restoring high-quality photorealistic faces from artistic portraits is challenging.

Motivated by such challenges, the recovery of photorealistic images from portraits has recently received some attention [34, 35, 10, 54]. The existing methods [34, 35, 10, 54] take a portrait image and then use an autoencoder to generate a photorealistic face image. These methods do not utilize the valuable semantic information in the process of face recovery. Despite training on large-scale datasets, they fail to provide consistent mappings between Stylized Portraits (SP) and ground-truth Real Faces (RF). Thus, they cannot preserve or enforce desired facial attributes in the recovered images. As shown in Fig. 1(d), 1(e) and 1(f), the facial details recovered by the state-of-the-art methods [34, 35, 10] are semantically and perceptually inconsistent with the ground-truth images. Inaccuracies range from an unnatural blur to attribute mismatches which include (but are not limited to) *Black Hair* and *Open Mouth*.

Unlike previous works, we propose to use facial attributes as high-level semantic information to boost the quality of recovered face images. Simply embedding the binary facial attribute vector as an additional input channel to the network results in visible distortions (see Fig. 4(e)). We observe that only low-frequency facial components are visible in the stylized input faces as a residual image (the difference between the RF image and the recovered face image) contains the missing high-frequency details. Thus, to recover the high-frequency facial details, we incorporate auxiliary facial attributes into the residual features.

Based on our observations above, we present a novel Face Recovery Network (FRN) that can use facial attributes during the recovery step. Our FRN uses an autoencoder with residual block-embedded skip connections to incorporate visual features obtained from portraits as well as semantic cues provided by facial attributes. FRN progressively upsamples the concatenated feature maps through its deconvolutional layers. Moreover, we employ a discriminative network that examines whether a recovered face image resembles an authentic face image and whether the attributes extracted from the recovered face are consistent with the input attributes. As a result, our discriminative

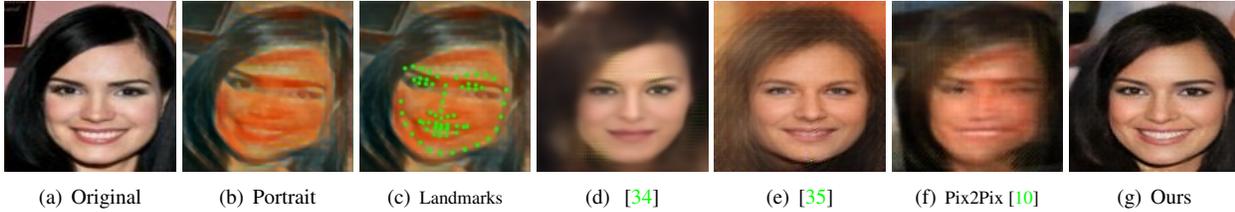


Figure 1. Comparisons to the state-of-the-art methods. (a) Ground-truth face image (from test dataset; not used in the training). (b) Unaligned stylized portraits of (a) from *Scream* style (unseen style in training), respectively. (c) Detected landmarks by approach [52]. (d) Results obtained by [34]. (e) Results obtained by [35]. (f) Results obtained by [10] (pix2pix). (g) Our results.

network can guide the generative network to incorporate semantic information into the recovery process. As shown in Fig. 1(g), our network learns consistent mappings between SP and RF facial patterns and preserves low-frequency details. Thus, we can generate realistic face images with details of the ground truth faces (e.g. *Black Hair, Smiling, Straight Hair, Wearing lip stick, pink cheeks*), as in Fig. 1(g).

We require a large number of pairs of Stylized Portraits (SP) and Real Face (RF) for training. Thus, we synthesize a large-scale training dataset. However, the choices of styles are numerous—we cannot generate all possible stylized faces for training. To select distinctive styles for training, we use a style-distance metric to measure the style distinctiveness. For this purpose, we use the Gram matrices [7] and the Log-Euclidean distance [12] although other non-Euclidean distances can also be explored [16, 15]. Specifically, we first measure the distance between Gram matrices of stylized images and the average Gram matrix of real faces, and then select the most distinctive styles, *i.e.* largest distance, for training. Furthermore, we note that our CNN filters learned from the data of seen styles (used for the training phase) can also extract informative features from images belonging to unseen styles. Thus, the facial information of unseen stylized portraits can be extracted and used to generate realistic faces, as later demonstrated in our experiments. The main contributions of our work can be summarized as follows:

- I. We design a novel framework which removes styles from unaligned stylized portraits. Our framework encodes stylized images with facial attributes and then recovers realistic faces from encoded feature maps.
- II. We propose an autoencoder with residual block-embedded skip-connections to extract residual feature maps from SP inputs and combine the extracted feature maps with facial attributes. In this fashion, we fuse visual and semantic information for best visual results.
- III. By manipulating input attribute vectors, our network can generate the realistic faces with desired attributes.

To the best of our knowledge, we are the first to use facial attributes for the face recovery from stylized portraits.

2. Related Work

Below we review papers related to neural style transfer and deep generative models for image generation.

2.1. Deep Generative Models

Recently, Generative Adversarial Networks (GANs) [9] have led to large improvements in image generation tasks. GANs learn the distribution of the training data in a non-conditional setting. Although these methods produce impressive photorealistic images, they cannot distinguish the identities of subjects. Recently, conditional GANs [10] were used to generate images conditioned on certain input variables. Conditional GANs benefit many applications such as super-resolution [47, 48, 17, 46], image generation [25, 14, 3, 51, 35], image inpainting [45, 27], general purpose image-to-image translation [10], image manipulation [53], synthesize faces from the landmarks [4], and style transfer [38]. In particular, Li and Wand [19] train a Markovian GAN for the style transfer via Markovian neural patches which capture local style statistics. Isola *et al.* [10] develop “pix2pix” framework which uses the patch-GAN to transfer low-level features from the input to the output domain. When these patch-based approaches are used to destylize portraits, they produce visual artefacts and fail to capture the global structure of the faces.

Sketch from/to photograph synthesis is explored in [24, 49, 37, 32, 30, 40]. When compared to sketch-to-face synthesis, viewed as a specific case of face recovery, our unified framework is able to process much more complex styles.

Recently, Yan *et al.* [44] used a conditional CNN to generate faces based on attributes. Perarnau *et al.* [28] developed an invertible conditional GAN to generate new faces by editing facial attributes of input images, while Shen and Liu [33] manipulated attributes of an input image via its residual image. As their methods are dedicated to generating new face images rather than the face recovery, they cannot preserve identity. In contrast, our method uses attributes to reduce the uncertainty of the face recovery and recover faithful realistic faces from artistic portraits.

2.2. Neural Style Transfer

Style transfer aims to synthesize an image with the visual content of the input image and a chosen style. The seminal work by [6] shows that correlation between fea-

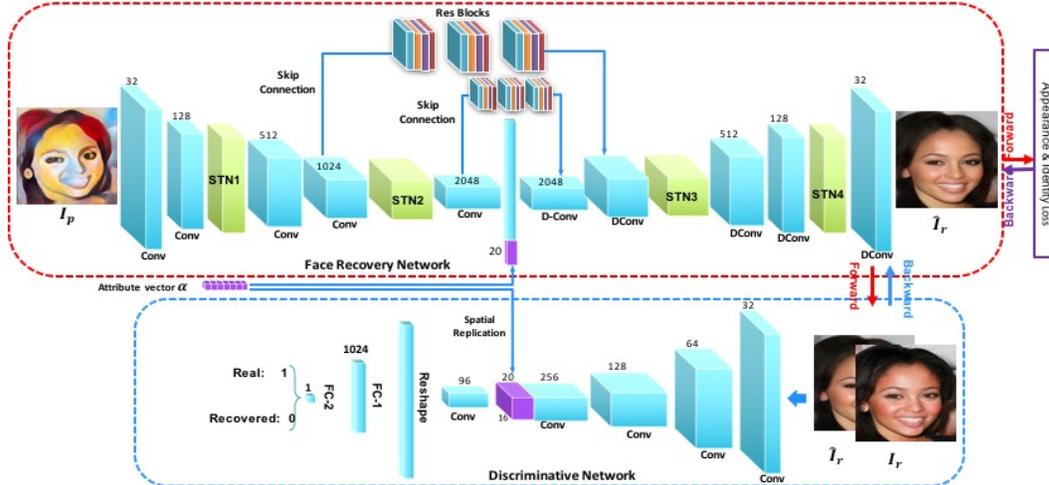


Figure 2. Our attribute-embedded face recovery framework has a generative network (red frame) and a discriminative network (blue frame).

ture maps (*i.e.*, Gram matrix) captures visual styles. Since then, many follow-up works used Gram-based objectives, such as iterative optimization methods [7, 8, 18, 43] and feed-forward networks [38, 13, 19] which are computationally costly due to the optimization step at the testing stage. In contrast, feed-forward methods learn a transformation to perform stylization in a feed-forward manner.

Johnson *et al.* [13] train a generative network for a fast style transfer via perceptual loss functions. Their generator network follows [29] and uses residual blocks. Texture Network [38] uses a multi-resolution generator network. Ulyanov *et al.* [39] replace the spatial batch norm. with the instance normalization for faster convergence. Wang *et al.* [41] enhance the granularity of the feed-forward style transfer via a multimodal CNN and stylisation losses.

These feed-forward methods are limited by the need to train one network per style due to the lack of generalization in network design. To deal with this restriction, recent approaches encode multiple styles within a single feed-forward network [5, 1, 20, 21]. Dumoulin *et al.* [5] use conditional instance normalization to learn necessary normalization parameters for each style. Given feature activations of the content and style images, [2] replaces content features with the style features patch-by-patch. To achieve an arbitrary style transfer, Chen *et al.* [1] swap content features with style features locally. Li *et al.* [20] adapt a single feed-forward network via a texture controller module.

As pointed by [35], direct use of neural style transfer for face recovery is suboptimal. Even though recent works [34, 35] are designed to destylize portrait images, they distort facial details and fail to recover facial traits (*e.g.*, hair color, lipstick, open/closed lips) to match the ground-truth. Since facial traits, such as hair color, are difficult to infer, the pixel-wise ℓ_2 norm and perceptual losses cannot yield correct facial attributes. Thus, state-of-the-art face destylization methods produce ambiguous results.

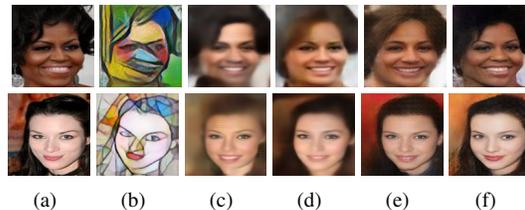


Figure 3. Contribution of each loss of AFRP network. (a) Ground-truth face images. (b) Input unaligned portraits from unseen styles. (c) Recovered faces without DN or identity-preserving loss. (d) Recovered faces with the ℓ_2 loss and discriminative loss. (e) Recovered faces with the ℓ_2 , discriminative and identity-preserving losses. (f) Our final results by embedding facial attributes.

3. Proposed Method

Below we present an attribute-guided framework for face recovery that takes SP images and facial attribute vectors as inputs, and it outputs photorealistic images of faces.

3.1. Network Architecture

Our network consists of two parts: a Face Recover Network (FRN) and a Discriminative Network (DN). FRN is composed of an autoencoder with skip connections with residual blocks. FRN extracts residual feature maps from input portraits and concatenates the corresponding 20-dimensional attribute vector with the extracted residual feature vector at the bottleneck of the autoencoder and then upsamples it. In this manner, we fuse visual and semantic information to attain high-quality visual performance. The role of DN is to guide the input attributes and the recovered face images to be similar to their real counterparts. The attribute vector is replicated and then concatenated with the extracted feature maps of the convolutional layer of DN. The entire architecture of our network is shown in Fig. 2.

FRN: This module employs a deep fully convolutional autoencoder for face recovery from portraits (see the red frame of Fig. 2). The convolutional layers of the encoder capture feature maps of input portraits while deconvolutional lay-

ers of the decoder upsample these feature maps to recover facial details. Previous works [34, 35, 10, 54] do not use valuable semantic information during face recovery. In contrast, our FRN incorporates low-level visual and high-level semantic information (*i.e.* facial attributes) for face recovery to reduce the ambiguity of mappings between SP and RF images. Specifically, at the bottleneck of the autoencoder, the attribute vector is concatenated with the residual feature vector (see the purple blocks in Fig. 2). A naive embedding of a semantic vector into SP may cause artefacts, *e.g.* Fig. 4(e) (encoding input portraits with attributes instead of residual feature maps) shows the identity is confused.

We also link symmetrically top convolutional and deconvolutional layers via skip-layer connections [23] as they pass higher-resolution visual details of portraits from convolutional to deconvolutional layers for better restoration quality. Each skip-connection comprises three residual blocks which help our network remove the styles of input portraits while increasing accuracy as in Fig. 4(g). Without skip-connections, we obtain blurred faces as in Fig. 4(c).

Note that input portraits are misaligned (in-plane rotations, translations). Similar to [35], we use multiple Spatial Transformer Networks (STNs) [11] in FRN (see the green blocks in Fig. 2) which compensate for misalignments of input portraits. Thus, our method does not require the use of facial landmarks or 3D face models.

To measure the appearance similarity between the recovered faces and their RF ground-truth counterparts, we use a pixel-wise ℓ_2 loss and an identity-preserving loss [35]. The pixel-wise ℓ_2 loss enforces intensity-based similarity between images of recovered faces and their ground-truth images. The autoencoder supervised by the ℓ_2 loss tends to output over-smoothed results as shown in Fig. 3(c). For the identity-preserving loss, we use FaceNet [31] to extract features from images (see Sec. 3.2 for more details), and then we compare the Euclidean distance between features of two images. In this way, we encourage feature similarity between the recovered faces and their ground-truth counterparts. Without the identity-preserving loss, the network produces random artefacts that resemble facial details, such as wrinkles, as shown in Fig. 3(d).

DN: In order to force the FRN to encode facial attribute information, we employ a conditional discriminative network. In particular, the discriminative network distinguishes whether the attributes of face images recovered by FRN match the desired attributes. DN also helps the recovered images to be similar to RF images. Since our FRN network may learn to ignore attribute vectors, *e.g.*, the weights corresponding to the semantic information are all zero, we design a discriminator network that incorporates semantic attribute information into the generative process. As shown in the first row of Fig. 4(f), the recovered hair color in the image is brown even if the ground truth hair color is black.

This implies that attribute cues are not exploited by the network. Thus, we design a discriminative network which uses attribute embedding in the learning process.

As shown in the blue frame of Fig. 2, DN consists of convolutional and fully connected layers. The real and recovered faces are fed into the network. The attribute information is fed into the middle layer of the network as a conditional information. As CNN filters in the first layers extract low-level features and filters in the higher layers extract semantically meaningful image patterns [50], in our experiment concatenating features maps with the attribute vectors in the fourth convolutional layer in DN yields better empirical results. When there is a mismatch between the extracted features and the input attributes, the discriminative network will pass the errors to the FRN network during backpropagation. As shown in Fig. 4(g), our final result matches the ground-truth facial expression, age and gender.

3.2. Training Procedure

To train AFRP end-to-end, we construct SP, RF and attribute vector triplets (I_p, I_r, \mathbf{a}) as our training dataset, where I_r is the aligned real face image, and I_p is the corresponding synthesized unaligned portrait image. For each RF, we synthesize different unaligned SP images from various artistic styles to obtain SP/RF training pairs. As detailed in Section 4, we only use stylized portraits from three distinct styles for training. We use SP image I_p and its ground-truth attribute label vector \mathbf{a} as inputs and the corresponding RF ground-truth image I_r as a target during training.

We train our FRN network using a pixel-wise ℓ_2 loss, a feature-wise loss and an adversarial loss to force the generated face I_r to resemble its corresponding ground-truth. In addition, we employ a binary cross-entropy loss to update our discriminative network. Since the STN layers are interwoven with the layers of our autoencoder, we optimize the parameters of the autoencoder and the STN layers simultaneously. Below we explain each loss individually.

Pixel-wise Intensity Similarity Loss: We train a feed-forward network to produce an aligned photorealistic face from a given unaligned portrait. To this end, we feed our FRN with I_p images and their corresponding attributes \mathbf{a} as inputs and then force the recovered face \hat{I}_r to be similar in intensity to its ground-truth counterpart I_r . Hence, we minimize the objective function \mathcal{L}_{pix} :

$$\begin{aligned} \mathcal{L}_{pix}(\Theta) &= \mathbb{E} \|\hat{I}_r - I_r\|_F^2 = \\ & \mathbb{E}_{(I_p, I_r, \mathbf{a}) \sim p(I_p, I_r, \mathbf{a})} \|\mathbf{G}_\Theta(I_p, \mathbf{a}) - I_r\|_F^2, \end{aligned} \quad (1)$$

where $\mathbf{G}_\Theta(I_p, \mathbf{a})$ and Θ represent the output and parameters of our FRN, respectively. We denote $p(I_p, I_r, \mathbf{a})$ as the joint distribution of the SP and RF images and the corresponding attributes in the training dataset.

Identity-preserving Loss: To obtain good identity-preserving results, we extract feature maps from the ReLU

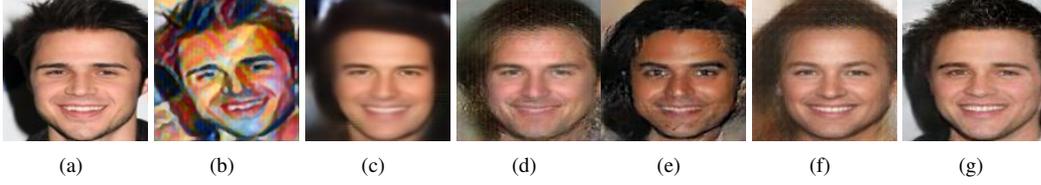


Figure 4. Ablation study. (a) RF ground-truth image. (b) Unaligned input portrait. (c) Result without skip connections/residual blocks in the autoencoder. (d) Result without residual blocks in the autoencoder. (e) Result when the attribute vector is concatenated with the SF input directly. (f) Result without attribute embedding. A standard discriminative network similar to the decoder in [35]. (g) Our final result.

activations of the FaceNet and compute the Euclidean distance between features of the recovered face $\hat{I}_r = G_{\Theta}(I_p, \mathbf{a})$ and ground-truth face I_r . As the FaceNet network is pre-trained on a large image dataset, it captures visually meaningful facial features and helps to preserve the identity information. Our identity-preserving loss \mathcal{L}_{id} is

$$\mathcal{L}_{id}(\Theta) = \mathbb{E} \|\psi(\hat{I}_r) - \psi(I_r)\|_F^2 = \mathbb{E}_{(I_p, I_r, \mathbf{a}) \sim p(I_s, I_r, \mathbf{a})} \|\psi(G_{\Theta}(I_p, \mathbf{a})) - \psi(I_r)\|_F^2, \quad (2)$$

where $\psi(\cdot)$ denotes the feature maps extracted from the layer ReLU3-2 of the FaceNet.

Discriminative Loss: The discriminative network should note if recovered faces contain the desired attributes and distinguish recovered faces from real ones. Moreover, FRN should make the discriminative network D_{Φ} fail to distinguish recovered faces from real ones and the attributes of generated faces should match the input attributes. Hence, in order to train the discriminative network, we take real FR face images I_r and their corresponding ground-truth attributes \mathbf{a} as positive sample pairs (I_r, \mathbf{a}) . Negative samples are constructed from recovered faces \hat{I}_r and their ground-truth attributes \mathbf{a} as well as real FR faces and mismatched (fake) attributes $\tilde{\mathbf{a}}$. Thus, the negative sample pairs consist of both (\hat{I}_r, \mathbf{a}) and $(I_r, \tilde{\mathbf{a}})$. The parameters of the discriminator Φ are updated by minimizing the loss:

$$\mathcal{L}_{dis}(\Phi) = -\mathbb{E}_{(I_r, \mathbf{a}) \sim p(I_r, \mathbf{a})} [\log D_{\Phi}(I_r, \mathbf{a})] - \mathbb{E}_{(\hat{I}_r, \mathbf{a}) \sim p(\hat{I}_r, \mathbf{a})} [\log(1 - D_{\Phi}(\hat{I}_r, \mathbf{a}))] - \mathbb{E}_{(I_r, \tilde{\mathbf{a}}) \sim p(I_r, \tilde{\mathbf{a}})} [\log(1 - D_{\Phi}(I_r, \tilde{\mathbf{a}}))], \quad (3)$$

where $p(I_r, \mathbf{a})$, $p(\hat{I}_r, \mathbf{a})$ and $p(I_r, \tilde{\mathbf{a}})$ denote distributions of real and recovered faces and the corresponding attributes respectively, and $p(I_r, \tilde{\mathbf{a}})$ represents the distribution of the recovered faces and the corresponding mismatched (fake) attributes. $D_{\Phi}(I_r, \mathbf{a})$, $D_{\Phi}(\hat{I}_r, \mathbf{a})$ and $D_{\Phi}(I_r, \tilde{\mathbf{a}})$ are the outputs of D_{Φ} . We update parameters of the discr. network, \mathcal{L}_{dis} loss is back-propagated to FRN.

Our FNR loss is a weighted sum of three terms: the pixel-wise loss, the discriminative loss, and the identity-preserving loss. The parameters Θ are obtained by minimizing the objective function of the FRN loss as follows:

$$\mathcal{L}_{FNR}(\Theta) = \mathbb{E}_{(I_p, I_r, \mathbf{a}) \sim p(I_p, I_r, \mathbf{a})} \|G_{\Theta}(I_p) - I_r\|_F^2 + \lambda \mathbb{E}_{I_p \sim p(I_p, \mathbf{a})} [\log D_{\Phi}(G_{\Theta}(I_p, \mathbf{a}))] + \eta \mathbb{E}_{(I_p, I_r, \mathbf{a}) \sim p(I_p, I_r, \mathbf{a})} \|\psi(G_{\Theta}(I_p, \mathbf{a})) - \psi(I_r)\|_F^2, \quad (4)$$

where λ determines a trade-off between the appearance and the attribute similarity, and η determines a trade-off between the image intensity and the feature similarity.

As $G_{\Theta}(\cdot)$ and $D_{\Phi}(\cdot)$ are differentiable, we apply back-propagation with respect to Θ and Φ , and optimize using Stochastic Gradient Descent (SGD) combined with the Root Mean Square Propagation (RMSprop).

3.3. Implementation Details

The discriminative network DN is needed in the training phase. In the testing phase, we take SP portraits and their corresponding attribute vectors as inputs and feed them to FRN. The outputs of FRN are the recovered photo-realistic face images. Although the attributes used for training are normalized between 0 and 1, they can be scaled up and down, e.g., above 1 or below 0, to manipulate the final results according to users' requirements.

We use convolutional layers with kernels of size 4×4 and stride 2 in the encoder and deconvolutional layers with kernels of size 4×4 and stride 2 in the decoder. We use mini-batches of size 64, the learning rate 10^{-3} and the decay rate 10^{-2} . For STNs, we use architectures as in [35]. In all experiments, λ and η are set to 10^{-2} and 10^{-3} , respectively, gradually reducing λ by a factor 0.995 to emphasize the importance of appearance similarity. As our method is feed-forward at the test time, it takes 8 ms to destylize a 128×128 image. Our code and model will be released on <http://claret.wikidot.com>.

4. Dataset and Preprocessing

To avoid overfitting in AFRP, a large number of SP/RF training image pairs are obtained from the CelebA dataset [22]. We randomly select 110K real faces for training and 2K images for testing. Then, we crop the central part of each image and resize it to 128×128 pixels as our RF ground-truth face images I_r . We augment RF images by rotation and translation. We use three distinct styles for synthesizing our training dataset (see Sec. 4.1). Finally, we obtain 330K SP/RF pairs and their corresponding attributes for training. We also use 2K unaligned real faces to synthesize 20K SP images from 10 diverse styles as our testing dataset. We also add sketches as an unseen style for testing. Training and testing datasets are disjoint. We choose 20 dominant attributes (*Bald, Bangs, Big nose, Black Hair, Blond Hair, Brown Hair, Eyeglasses, Gray Hair, Heavy Makeupe, Male,*

Table 1. Impact of tuning attributes on the classification results.

Attributes	GT Attr.	Increased	Decreased
	Acc.	Attr. Acc.	Attr. Acc.
<i>Young</i>	95%	100%	0.5%
<i>Male</i>	100%	100%	1%
<i>Beard</i>	79%	100%	15%

Mouth Open, Mustache, Narrow Eyes, No Beard, Pale Skin, Smiling, Straight Hair, Wavy Hair, Wearing Lipstick and Young) from 40 attributes in CelebA. The ground truth attributes are binary 0/1 values.

4.1. Style Distance Metric

It is not practical to generate a large number of styles for training. Thus, we propose a style distance metric to select the most difficult styles for the face recovery process. To this end, we compute Gram matrices for various styles from feature maps of pre-trained VGG-network [36]. Then, we measure the similarity of styles via the Log-Euclidean metric [12] between Gram matrices of style images and the average Gram matrix of real training face images. As a result, we choose *Candy, Wave* and *Mosaic* styles for training.

5. Experiments

We compare our approach qualitatively/quantitatively to the state-of-the-art methods [13, 34, 10, 55, 35]. For fairness, we retrain these methods on our training dataset for the task of photorealistic face recovery from stylized portraits.

5.1. Attribute Manipulation in Face Recovery

By manipulating attribute vectors, we can also post-edit results. As shown in Fig. 5(f), by changing the hair color attribute, we can restore the same person with different hair color. Our method can also manipulate the age in Fig. 5(b), remove the eye-lines and lipstick in Fig. 5(c), open or close mouths in Fig. 5(d), add beard in Fig. 5(e), and change the hair color in Fig. 5(f).

To test if the attribute information has been successfully embedded in our network, we choose three different attributes, *i.e.* *Young, Male* and *Beard*, and we train an attribute classifier for each attribute. By increasing and decreasing the corresponding attribute values, the true positive accuracies change accordingly, as shown in Table 1. This indicates that the attribute information and thus semantic information has been successfully injected into recovery.

5.2. Qualitative Evaluation

We provide sample results in Fig. 6 (also see our supp. material). Note that [34, 10, 13, 55] require input SP faces to be aligned before recovery so we employ an STN to align all the SP images. Our method and [35] automatically generate upright real face images. The aligned upright RF ground-truth images are shown for comparison. We visually compare our approach with five methods detailed below.

Johnson *et al.* [13] captures correlation between feature maps of the portrait and the synthesized face (Gram matrices).

We retrain this approach for destylization. However, it fails to preserve spatial structures of face images and network generates distorted facial details and unnatural artefacts. As shown in Fig. 6(c), the facial details are blurred and the artistic styles have been removed only partially.

Shiri *et al.* [34] introduce a face destylization which uses only a pixel-wise loss in their generative network and a standard discriminator to enhance facial details. Although their approach is trained on a large-scale dataset, it fails to generate authentic facial details due to the existence of various styles. As seen in Fig. 6(d), it produces distorted results and the facial colors are inconsistent. It cannot recover faces from unaligned portraits or large pose portraits.

Isola *et al.* [10] train a "U-net" generator augmented with a PatchGAN discriminator in an adversarial framework, known as "pix2pix". Their network does not capture the global structure of faces. As shown in Fig. 6(e), pix2pix can generate acceptable results for the seen styles but it fails to remove the unseen styles and produces obvious artefacts.

CycleGAN [55] is an image-to-image translation method that uses unpaired datasets. Since CycleGAN also employs a patch-based discriminator, it cannot capture the global structure of faces either. As CycleGAN uses unpaired face datasets, the low-level features of the stylized faces and real faces do not match well. As shown in Fig. 6(f), this method produces distorted results and does not preserve the identities with respect to the input images.

Shiri *et al.* [35] exploit an identity-preserving loss to reveal the photorealistic faces from unaligned stylized faces. They also employ a simple autoencoder and standard discriminative network to recover the real faces, but their discriminative network is only used to force the generative network to produce sharper results without imposing attribute information. As shown in Fig. 6(g), their method suffers mismatched hair colors. As shown in the third row of Fig. 6(g), their method also recovers male facial details.

In contrast, our results demonstrate higher fidelity and better consistency with respect to the ground-truth face images as shown in Fig. 6(h). Our method, evaluated on portraits from seen/unseen styles and sketches, produces high-quality realistic faces which also match the semantic composition of ground-truth images. Our network recovers the photorealistic faces from various stylized portraits of the same person as shown in Fig. 6. Note that the recovered faces resemble each other. This demonstrates the robustness of our network with respect to different styles.

5.3. Quantitative Evaluation

Face Reconstruction Analysis. To evaluate the reconstruction performance, we measure the average Peak Signal to Noise Ratio (PSNR) and Structural Similarity (SSIM) [42] on the entire test dataset. Table 2 indicates that our method achieves superior quantitative performance in comparison to other methods on both seen and unseen

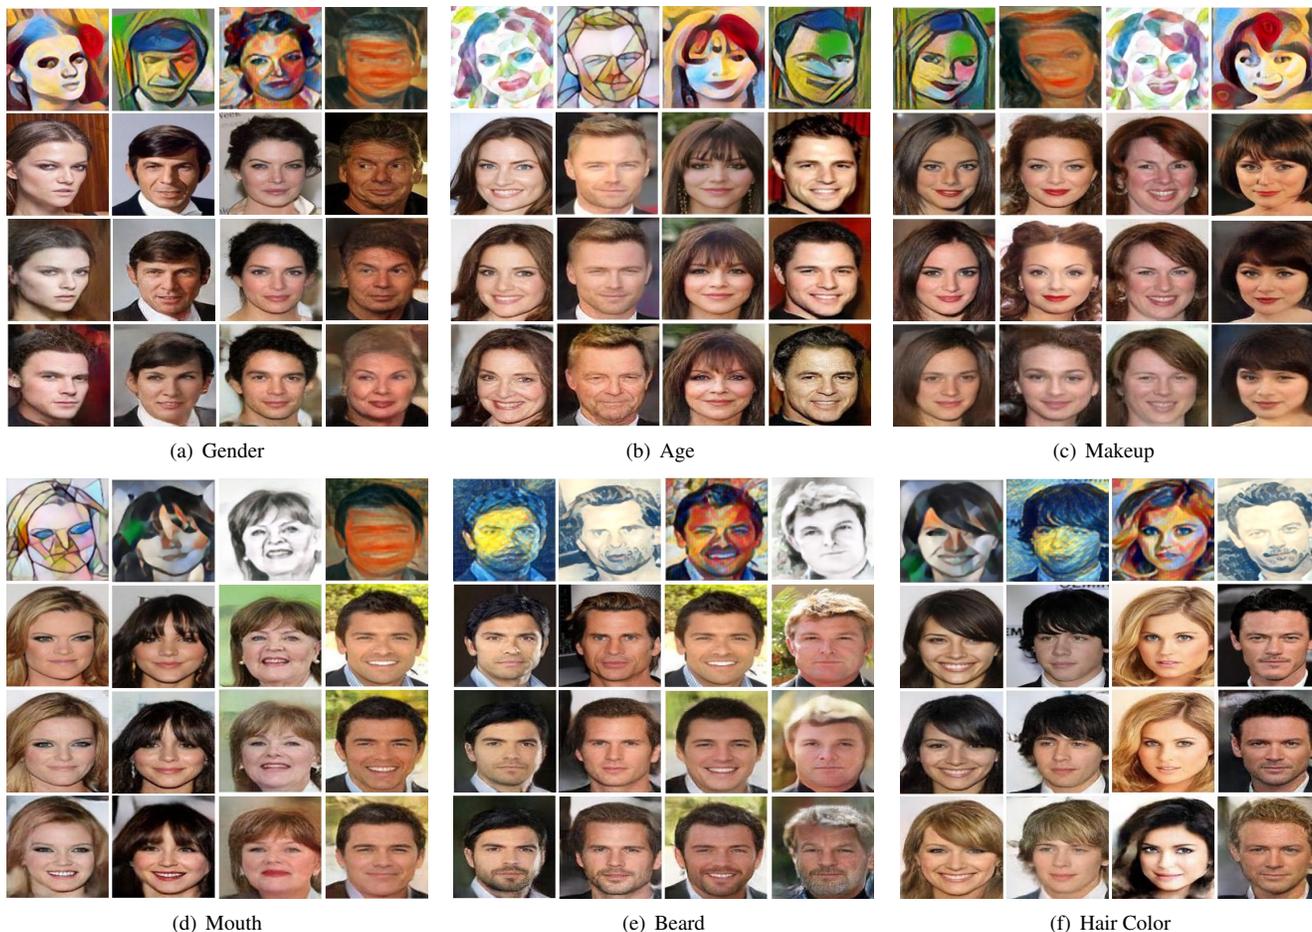


Figure 5. Our method lets us fine-tune the recovered results by manipulating the attributes. First row: Unaligned input portraits. Second row: RF ground-truth faces. Third row: Our results with ground-truth attributes. Fourth row: Our results by adjusting attributes. (a) Changing gender. (b) Adding age. (c) Removing makeup. (d) Opening and closing mouth. (e) Adding beard. (d) Changing hair color.

styles. As indicated in Table 2, we show the quantitative results of solely using FRN, marked as FRN. Also, the results of using both FRN and a standard DN, indicated by FRN+SDN, is demonstrated in Table 2. The standard DN only forces FRN to generate realistic faces, and thus it improves the results qualitatively and quantitatively. Since FRN augmented with attributes may learn a trivial solution, where all attribute vectors will be neglected, using a standard DN cannot force FRN to embed such attribute information. On the contrary, our conditional DN is able to distinguish whether the attributes match the input faces or not, thus forcing FRN to embed attribute information in the process of face recovery. Thus, the ambiguity is significantly reduced and the network achieves better performance.

Face Retrieval Analysis. To demonstrate that the faces recovered by our method are highly consistent with their ground-truth counterparts, we run a face recognition algorithm [26] on our test dataset for both seen and unseen styles. For each investigated method, we consider 2K recovered faces from one style as query images and then use their ground-truth real faces as a gallery dataset. We run [26]

Table 2. Comparisons of PSNR and SSIM on the entire test dataset.

Method	Seen Styles		Unseen Styles		Unseen Sketches	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Johnson [13]	17.85	0.76	18.07	0.75	18.11	0.76
Shiri [34]	19.22	0.81	19.09	0.80	19.01	0.80
pix2pix [10]	18.45	0.78	18.12	0.77	18.07	0.78
CycleGAN [55]	18.35	0.75	18.29	0.75	18.08	0.76
Shiri [35]	19.35	0.80	19.31	0.79	19.045	0.80
FRN	18.42	0.76	18.51	0.75	18.49	0.75
FRN + SDN	19.66	0.81	19.58	0.80	19.62	0.80
AFRP	20.01	0.84	19.99	0.83	19.98	0.83

to check whether the correct person is retrieved within the top-5 matched images and then an average retrieval score is obtained. We repeat this procedure for each style and then obtain the average Face Retrieval Ratio (FRR) by averaging all scores from the seen and unseen styles, respectively. As indicated in Table 3, our AFRP network outperforms the other methods across all the styles. Even for the unseen styles, our method can still generate realistic facial details with high fidelity to the ground-truth.

5.4. Destylizing Original Paintings and Sketches

Fig. 7 illustrates that our method is not limited to computer-generated stylized portraits and it can also effi-

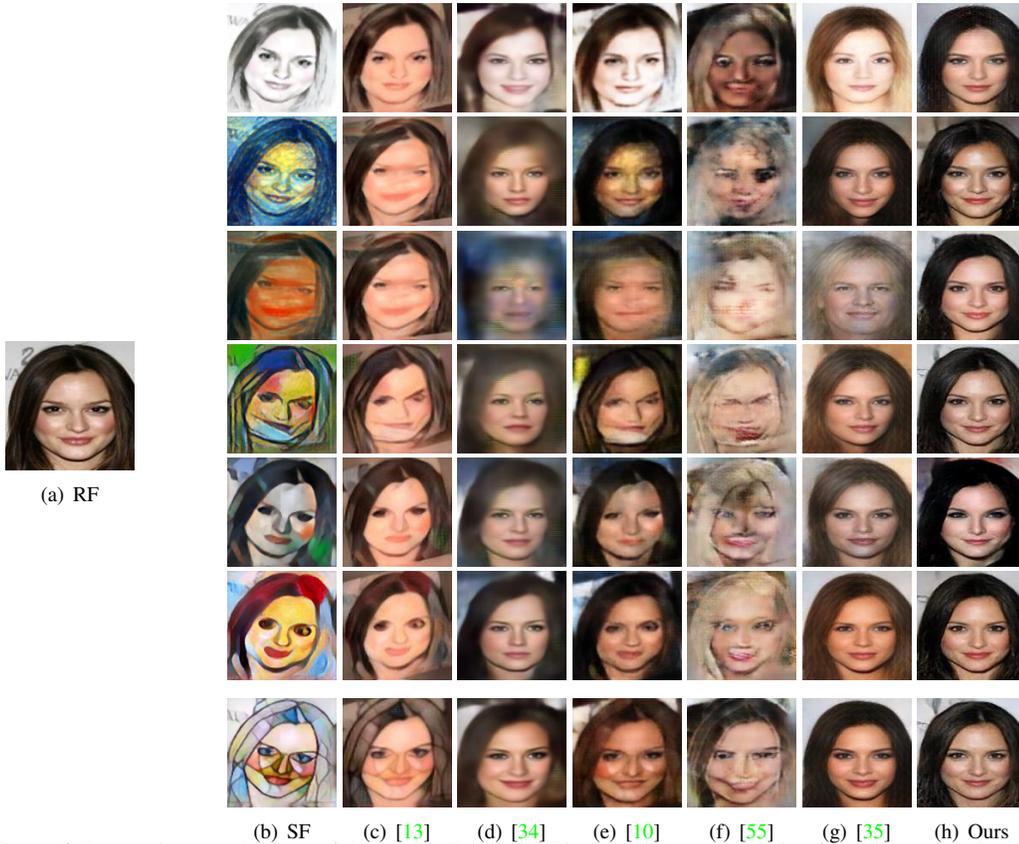


Figure 6. Comparisons to the state of the art. (a) Original RF image. (b) Input portraits (from the test set) including the unseen styles *Sketch*, *Starry*, *Scream*, *La Muse* and *Udnie* and the seen styles *Candy* and *Mosaic*. (c) Johnson *et al.*'s method [13]. (d) Shiri *et al.*'s method [34] (e) Isola *et al.*'s method [10] (pix2pix). (f) Zhu *et al.*'s method [55] (CycleGAN). (g) Shiri *et al.*'s method [35]. (h) Our method.

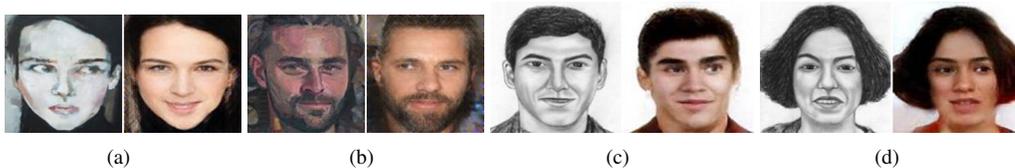


Figure 7. Results for the original unaligned paintings and hand-drawn sketches. Right: the original portraits. Left: our results.

ciently recover photorealistic faces from original paintings and sketches. We choose real paintings from art galleries and hand-drawn sketches as our test examples. Since we do not know the ground-truth attributes, we set the attribute vectors to neutral values *i.e.*, 0.5. As shown in Fig. 7, despite attributes may be inaccurate, our method still generates authentic face images regardless of their original styles.

6. Conclusions

We have introduced an attribute guided generative-discriminative network to recover photorealistic faces from unaligned stylized portraits in an end-to-end fashion. With the help of the conditional discriminative network, our network successfully incorporates facial attribute vectors into the residual features of input portraits at the bottleneck of the autoencoder. Our network is able to preserve the iden-

Table 3. Comparisons of FRR on the entire test dataset.

Method	Seen Styles	Unseen Styles	Unseen Sketch
Johnson [13]	55.57%	50.48%	54.36%
Shiri [34]	78.00%	66.89%	65.26%
pix2pix [10]	76.03%	62.67%	64.64%
CycleGAN [55]	36.07%	33.68%	32.75%
Shiri [35]	84.51%	75.32%	75.44%
AFRP	93.08%	83.14%	92.05%

tity of generated faces and it can post-edit the recovered results by adjusting the attribute information. Moreover, our algorithm demonstrates good generalization ability for recovery of portraits from unseen styles, real paintings as well as hand-drawn sketches.

Acknowledgement. This work is supported by the Australian Research Council (ARC) grant DP150104645.

References

- [1] D. Chen, L. Yuan, J. Liao, N. Yu, and G. Hua. Stylebank: An explicit representation for neural image style transfer. *arXiv preprint arXiv:1703.09210*, 2017. 3
- [2] T. Q. Chen and M. Schmidt. Fast patch-based style transfer of arbitrary style. *arXiv preprint arXiv:1612.04337*, 2016. 3
- [3] E. L. Denton, S. Chintala, R. Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *NIPS*, 2015. 2
- [4] X. Di, V. A. Sindagi, and V. M. Patel. Gp-gan: gender preserving gan for synthesizing faces from landmarks. *arXiv preprint arXiv:1710.00962*, 2017. 2
- [5] V. Dumoulin, J. Shlens, and M. Kudlur. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629*, 2016. 3
- [6] L. Gatys, A. S. Ecker, and M. Bethge. Texture synthesis using convolutional neural networks. In *NIPS*, 2015. 2
- [7] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016. 2, 3
- [8] L. A. Gatys, A. S. Ecker, M. Bethge, A. Hertzmann, and E. Shechtman. Controlling perceptual factors in neural style transfer. In *CVPR*, 2017. 3
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [10] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint arXiv:1611.07004*, 2016. 1, 2, 4, 6, 7, 8
- [11] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *NIPS*, 2015. 4
- [12] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi. Kernel methods on the riemannian manifold of symmetric positive definite matrices. In *CVPR*, 2013. 2, 6
- [13] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 3, 6, 7, 8
- [14] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *stat*, 2014. 2
- [15] P. Koniusz, Y. Tas, H. Zhang, M. Harandi, F. Porikli, and R. Zhang. Museum exhibit identification challenge for the supervised domain adaptation and beyond. *ECCV*, pages 788–804, 2018. 2
- [16] P. Koniusz, H. Zhang, and F. Porikli. A deeper look at power normalizations. *CVPR*, pages 5774–5783, 2018. 2
- [17] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint*, 2017. 2
- [18] C. Li and M. Wand. Combining markov random fields and convolutional neural networks for image synthesis. In *CVPR*, 2016. 3
- [19] C. Li and M. Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *ECCV*, 2016. 2, 3
- [20] Y. Li, C. Fang, J. Yang, Z. Wang, X. Lu, and M.-H. Yang. Diversified texture synthesis with feed-forward networks. *arXiv preprint arXiv:1703.01664*, 2017. 3
- [21] Y. Li, C. Fang, J. Yang, Z. Wang, X. Lu, and M.-H. Yang. Universal style transfer via feature transforms. In *NIPS*, 2017. 3
- [22] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 5
- [23] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 4
- [24] H. Nejadi and T. Sim. A study on recognizing non-artistic face sketches. In *WACV*, 2011. 2
- [25] A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016. 2
- [26] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *BMVC*, 2015. 7
- [27] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 2
- [28] G. Perarnau, J. van de Weijer, B. Raducanu, and J. M. Álvarez. Invertible conditional gans for image editing. *arXiv preprint arXiv:1611.06355*, 2016. 2
- [29] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 3
- [30] P. Sangkloy, J. Lu, C. Fang, F. Yu, and J. Hays. Scribbler: Controlling deep image synthesis with sketch and color. In *CVPR*, 2017. 2
- [31] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015. 4
- [32] A. Sharma and D. W. Jacobs. Bypassing synthesis: Pls for face recognition with pose, low-resolution and sketch. In *CVPR*, 2011. 2
- [33] W. Shen and R. Liu. Learning residual images for face attribute manipulation. *arXiv preprint arXiv:1612.05363*, 2016. 2
- [34] F. Shiri, X. Yu, P. Koniusz, and F. Porikli. Face destylization. In *DICTA*. IEEE, 2017. 1, 2, 3, 4, 6, 7, 8
- [35] F. Shiri, X. Yu, F. Porikli, R. Hartley, and P. Koniusz. Identity-preserving face recovery from portraits. *WACV*, 2018. 1, 2, 3, 4, 5, 6, 7, 8
- [36] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [37] X. Tang and X. Wang. Face sketch synthesis and recognition. In *ICCV*, 2003. 2
- [38] D. Ulyanov, V. Lebedev, A. Vedaldi, and V. S. Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *ICML*, 2016. 2, 3
- [39] D. Ulyanov, A. Vedaldi, and V. Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *CVPR*, 2017. 3
- [40] L. Wang, V. Sindagi, and V. Patel. High-quality facial photo-sketch synthesis using multi-adversarial networks. In *Automatic Face & Gesture Recognition (FG 2018)*, 2018 13th IEEE International Conference on, pages 83–90. IEEE, 2018. 2

- [41] X. Wang, G. Oxholm, D. Zhang, and Y.-F. Wang. Multi-modal transfer: A hierarchical deep convolutional neural network for fast artistic style transfer. *arXiv preprint*, 2017. 3
- [42] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004. 6
- [43] P. Wilmot, E. Risser, and C. Barnes. Stable and controllable neural texture synthesis and style transfer using histogram losses. *arXiv preprint arXiv:1701.08893*, 2017. 3
- [44] X. Yan, J. Yang, K. Sohn, and H. Lee. Attribute2image: Conditional image generation from visual attributes. In *ECCV*, 2016. 2
- [45] R. Yeh, C. Chen, T. Y. Lim, M. Hasegawa-Johnson, and M. N. Do. Semantic image inpainting with perceptual and contextual losses. *arXiv preprint arXiv:1607.07539*, 2016. 2
- [46] X. Yu, B. Fernando, B. Ghanem, F. Porikli, and R. Hartley. Face super-resolution guided by facial component heatmaps. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 217–233, 2018. 2
- [47] X. Yu and F. Porikli. Face hallucination with tiny unaligned images by transformative discriminative neural networks. In *AAAI*, 2017. 2
- [48] X. Yu and F. Porikli. Hallucinating very low-resolution unaligned and noisy face images by transformative discriminative autoencoders. In *CVPR*, 2017. 2
- [49] P. C. Yuen and C. Man. Human face image searching system using sketches. *IEEE SMC*, 2007. 2
- [50] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014. 4
- [51] H. Zhang, V. Sindagi, and V. M. Patel. Image de-raining using a conditional generative adversarial network. *arXiv preprint arXiv:1701.05957*, 2017. 2
- [52] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Facial landmark detection by deep multi-task learning. In *ECCV*, 2014. 2
- [53] J.-Y. Zhu, P. Krähenbühl, E. Shechtman, and A. A. Efros. Generative visual manipulation on the natural image manifold. In *ECCV*, 2016. 2
- [54] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint arXiv:1703.10593*, 2017. 1, 4
- [55] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint arXiv:1703.10593*, 2017. 6, 7, 8