Two-phase Hair Image Synthesis by Self-Enhancing Generative Model

Haonan Qiu[12] , **Chuan Wang**[3] , **Hang Zhu**[1] , **Xiangyu Zhu**[1] , **Jinjin Gu**[1] , **Xiaoguang Han**[12*]

[1]The Chinese University of Hong Kong, Shenzhen
[2]Shenzhen Research Institute of Big Data, CUHK(SZ)   [3]Megvii (Face++) USA

{haonanqiu,hangzhu,jinjingu}@link.cuhk.edu.cn, uhzoaix@gmail.com
wangchuan@megvii.com, hanxiaoguang@cuhk.edu.cn

Abstract

Generating plausible hair image given limited guidance, such as sparse sketches or low-resolution image, has been made possible with the rise of Generative Adversarial Networks (GANs). Traditional image-to-image translation networks can generate recognizable results, but finer textures are usually lost and blur artifacts commonly exist. In this paper, we propose a two-phase generative model for high-quality hair image synthesis. The two-phase pipeline first generates a coarse image by an existing image translation model, then applies a re-generating network with self-enhancing capability to the coarse image. The self-enhancing capability is achieved by a proposed structure extraction layer, which extracts the texture and orientation map from a hair image. Extensive experiments on two tasks, Sketch2Hair and Hair Super-Resolution, demonstrate that our approach is able to synthesize plausible hair image with finer details, and outperforms the state-of-the-art.

1 Introduction

Accompanied with the success of applying conditional Generative Adversarial Networks (cGANs) [Mirza and Osindero, 2014] on image to image translation tasks [Isola *et al.*, 2017], generating realistic photos from sparse inputs, such as label maps [Chen and Koltun, 2017; Lassner *et al.*, 2017] and sketches [Portenier *et al.*, 2018], nowadays draws much attention of the researchers in both computer graphics and computer vision communities. Portrait image generation, as one of the most popular topics among generative tasks, has been widely studied [Wang *et al.*, 2018b; Karras *et al.*, 2017]. Although so, the hair regions, as one of most salient areas, are usually generated with blurry appearances.

To touch the heart of this problem, in this paper, we explore approaches to produce realistic hair photographs conditioned on sparse/low-resolution inputs, such as a hair sketch or a downsampled hair image. Hair is very different from other categories of images as its special textures: it is rendered from

(a) Hand-drawn sketch   (b) Coarse result by pix2pix

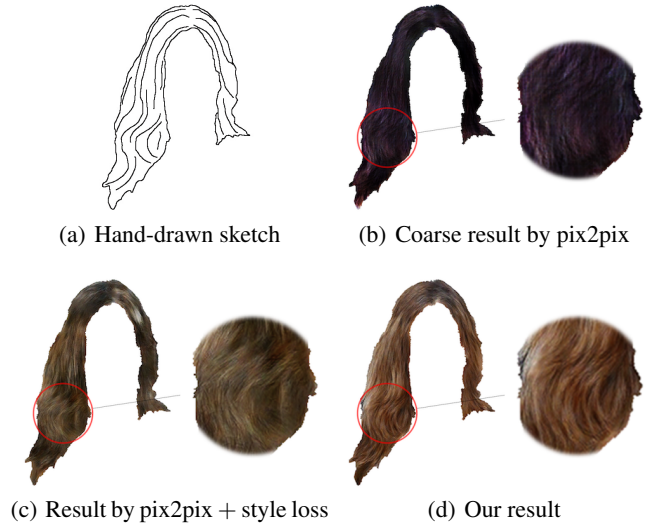(c) Result by pix2pix + style loss   (d) Our result

Figure 1: Given an image with limited guidance, such as a hand-drawn sketch (a), the basic pix2pix framework causes blur artifacts (b). Involving a style loss in training the pix2pix network helps to generate the structure with limited power (c). (d) Results by our two-phase approach, with finer strands and smoother texture synthesized.

thousands of long yet thin strands, being full of textured details. This makes the existing cGAN-based methods, such as pix2pix [Isola *et al.*, 2017], fail from two aspects: 1) its discriminator, by encoding the output image into a latent space, guarantees realistic output in a global fashion yet lacks constraints on local sharpness; 2) the input, as the condition of cGANs, is too weak to generate strand-wise pixels.

Based on our experiments, the first issue can be addressed by borrowing feature matching loss from [Wang *et al.*, 2018b] and style loss from [Gatys *et al.*, 2016] to guide the learning procedure. More importantly, to overcome the second challenge and make strands floating out, a self-enhancing generative model is proposed. Our key idea is conducting the generation into two phases: 1) we first utilize the state-of-the-art methods to produce a coarse-level output; 2) the strand-ware structures, such as orientation, is thus extracted. This is then treated as an enhanced condition and fed into a re-generation network, also based on cGANs. To support an end-to-end
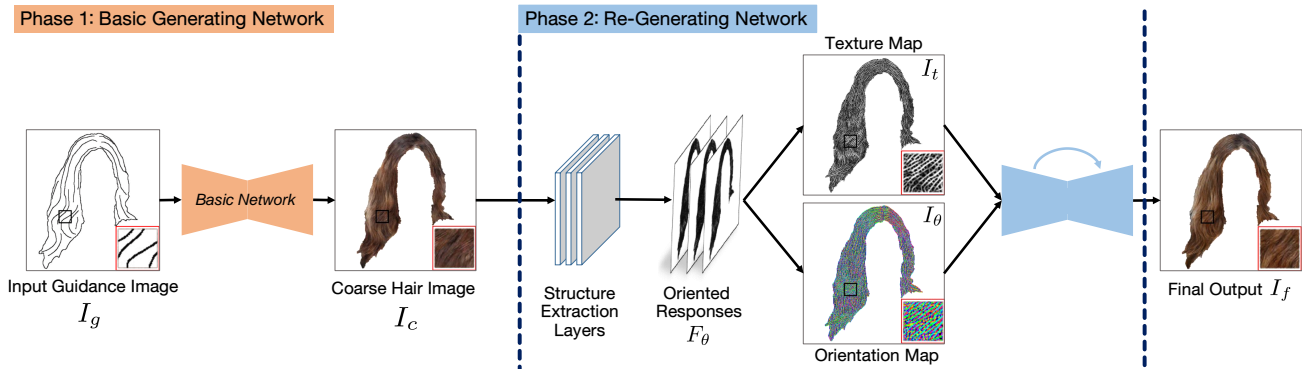
---

*Corresponding Author

Figure 2: The architecture of our framework is composed of two phases. In Phase 1, the input image with limited guidance I_g is transformed to a coarse image I_c with a basic net G_b. In Phase 2, a re-generating network G_r with self-enhancing capability is attached at the end of G_b, and produces the final output I_f. Self-enhancement is achieved by our proposed structure extraction layer, extracting texture and orientation map $[I_t, I_\theta]$ as the inputs for a U-net structure. The specific structures of I_g, G_b and G_r is task-specific. Here we use the task of Sketch2Hair to illustrate, while it is similar for the task of Hair Super-Resolution.

training of such two-phase network, a novel differentiable texture-extraction layer is embedded in the re-generating network, which enables its capability of self-enhancement.

To validate the effectiveness of the proposed self-enhancing mechanism, two tasks are studied: sketch2hair aims to realistic hair generation from a sparse sketch input; hair super-resolution targets generating high-resolution appearances from a down-sampled image. A high-quality dataset is also built to support conducting these two applications. Both the user study and visual comparisons show the superiorities of the proposed method against all existing ones. Thanks to the proposed structure extraction layer, all of the coarse hair image can be significantly enhanced.

In summary, our main contributions include:

- A novel self-enhancing module is designed, with which, the hair generation is modeled as an end-to-end two-phase framework. As demonstrated in two tasks, Sketch2Hair and Hair Super-Resolution, this strategy effectively benefits the state-of-the-art generative models for hair synthesis. We foresee this general strategy could be potentially to applied more hair synthesis related tasks.

- For the task of Sketch2Hair, we are the first application targeting realistic hair image synthesis purely from sketches. This can be regarded as a prototype for real-time sketch-based hairstyle editing.

- We constructed a high-quality hair dataset including 640 high-resolution hair photos with their corresponding sketches manually drawn. This database will be released upon our paper is accepted to facilitate the research in this field.

2 Related Work

Realistic hair synthesis. Generating virtual hairstyles is a long-standing research topic in computer graphics due to the important role it plays in representing human characters in games and movies. Most previous works focus on producing 3D hairs, according to user interactions [Mao *et al.*, 2004;

Fu *et al.*, 2007; Hu *et al.*, 2015] or real-captured images [Wei *et al.*, 2005; Jakob *et al.*, 2009; Chai *et al.*, 2016; Zhou *et al.*, 2018]. Given images, thanks to these modeling techniques, the hair can be recovered strand by strand which enables intelligent hair editing [Chai *et al.*, 2012; Chai *et al.*, 2013] or interpolation [Weng *et al.*, 2013] by performing manipulation in 3D space and then being re-rendered to 2D domain. Although these methods are able to result in realistic appearances, high computational cost are incurred due to the involvement of 3D matters. To avoid high computational cost in hair rendering, [Wei *et al.*, 2018] proposes a deep learning based hair synthesis method, which can generate high-quality results from an edge activation map. However, to obtain the activation map, an input CG hair model is still required for the initial rendering. In comparison, our method involves no 3D rendering module, and relies on a 2D image with sparse information only. With such a limited input, we still synthesize photo-realistic results thanks to the self-enhancing module as proposed.

In contrast, [Chen and Zhu, 2006] put forward a 2D generative sketch model for both hair analysis and synthesis, with which, a hair image can be encoded into a sketch graph, a vector field and a high-frequency band. Taking an image as input, such a multi-level representation provides a straightforward way to synthesize new hairstyles by directly manipulating the sketch. Compared with this work, our approach is feasible to infer a realistic output by only taking in a sparse sketch without any reference photo. We use deep neural networks to achieve spatial consistent conversion from sketches to colorful images instead of traditional image render, which is relatively time-consuming. To our knowledge, we are the first work using cGANs aiming at sketch to hair image synthesis.

Portrait super-resolution. Dong et al. [Dong *et al.*, 2014; Dong *et al.*, 2016] pioneered the use of convolutional networks for image super-resolution, achieving superior performance to previous works. Since only three layers of simple convolutional networks are used, SRCNN is still less effective in recovering image details. Later many deeper

and more effective networks structures [Kim *et al.*, 2016; Ledig *et al.*, 2017; Lim *et al.*, 2017; Zhang *et al.*, 2018; Wang *et al.*, 2018c] are designed and achieve great success in improving the quality of recovered images. As a hot object, several CNN based architectures [Zhu *et al.*, 2016; Cao *et al.*, 2017; Huang *et al.*, 2017] have been specifically developed for face hallucination. Also vectorization can also be utilized for image super-resolution as in [Lai *et al.*, 2009; Wang *et al.*, 2017]. To produce more details in results, adversarial learning [Yu and Porikli, 2016; Yu and Porikli, 2017] is also introduced. Recently, Li et al. [Li *et al.*, 2018] propose a semi-parameter approach to reconstruct high quality face image from unknown degraded observation with the help of reference image. However, as an important part of the portrait that fits closely to the face, recovered hair results are always neglected. The hair area recovered by current advanced approaches is always blurred or gelatinous, becoming a short board in portrait Super-Resolution. In this paper, as hair area segmentation is achieved by [Levinshtein *et al.*, 2018], we propose an extra hair textured enhancement to improve the recovered hair visual-quality. Our enhanced structure is able to be attached to almost super-resolution methods in an end-to-end manner, which can be regarded as an additional texture enhancement module.

Enhancing technology in generation. In generative task, spatial consistent conversion between two domains is hard to train especially when the transfer gap is large. Several enhanced methods from different direction to reduce transfer by dividing the whole conversion into some subtasks. To stabilize the training process, [Chen and Koltun, 2017; Wang *et al.*, 2018b; Karras *et al.*, 2017] propose the enhanced method that train the network from small scale and then fine tune on large scale. Also, in semantic segmentation, the strategy of regenerating from preliminary predictions to improve accuracy is widely used [Shelhamer *et al.*, 2017; Zheng *et al.*, 2015].

In low-level vision task, to recover the structure in generated result, [Xu *et al.*, 2015; Liu *et al.*, 2016] develop networks to approximate a number of filters for edge-preserving. Furthermore, [Pan *et al.*, 2018] proposes DualCNN which consists of two parallel branches to recover the structure and details in an end-to-end manner respectively. But in hair generation, generated texture, which can be regarded as structure, is prone to be messed up with some disordered noise. To overcome this problem, we extract structure from the coarse result and use it to regenerate our final result in an end-to-end manner. As far as we know, this enhanced strategy in CNN based generative model has not been proposed before. And our experimental results show this enhanced structure has significant effects to make less blurry area and more meticulous textures.

3 Network Architecture

Our proposed generative model for hair image synthesis takes an image of limited guidance I_g as input and produces a plausible hair image I_f as output. It enables the high-quality generation by a two-phase scheme. In Phase 1, a basic network $G_b(\cdot)$ as pix2pix in [Isola *et al.*, 2017] is used to gen-

(a) $\theta = 0$ (b) $\theta = \frac{\pi}{4}$ (c) $\theta = \frac{\pi}{2}$ (d) $\theta = \frac{3\pi}{4}$
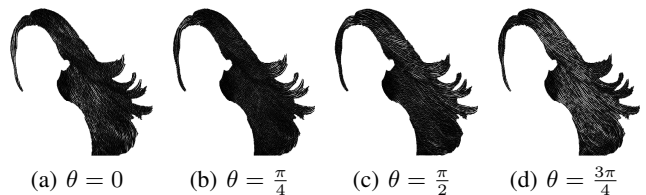
Figure 3: Gabor responses in different orientations. Hairline has a large response for its corresponding oriented Gabor filter. With this property, the core structure of hair can be extracted by the designed Gabor filter bank.

erate a coarse result I_c, which usually contains little texture and some blur artifacts. Then a re-generating network $G_r(\cdot)$ with self-enhancing capability is applied to produce the high-quality result I_f in Phase 2. We illustrate the entire framework in Figure 2.

3.1 Basic Generating Network

Given a hair-related image of limited guidance I_g, in Phase 1, we conduct an image-to-image translation between I_g and the target domain, i.e. a plausible hair image. The network is commonly a conditional GAN as in [Radford *et al.*, 2015; Isola *et al.*, 2017]. Specifically, in the tasks of Sketch2Hair (S2H) and Hair Super-Resolution (HSR), I_g are a sparse sketch image (Figure 1(a)) and low-resolution image (Figure 10(a)) respectively. Accordingly, the target image I_c is of the same resolution as I_g for S2H, and of a user-specified higher resolution for HSR. The network structures are also task specified. For example, in HSR the structure contains several more upsampling layers compared with S2H. We refer readers to [Isola *et al.*, 2017] for detailed descriptions of image-to-image translation networks for the two tasks. Also, for simplicity, we use "basic net" or "basic network" for an abbreviations of "basic generating network" in this paper.

3.2 Self-Enhancing Re-Generating Network

The hair image produced by Phase 1, I_c, is usually recognisable with core structure and close to the target. However, it is still far from plausible due to its lack of gloss, texture and strand-like style. To generate a high-quality hair image I_f, we further feed I_c into a re-generating network $G_r(\cdot)$ with self-enhancing capability, which is achieved by a newly-introduced *structure extraction layer* as follows.

Structure Extraction Layer

Given a hair image I, we follow the pipeline in [Chai *et al.*, 2012] to filter I by a set of oriented filters $\{K_\theta\}$, generating a set of oriented responses $\{F_\theta\}$, i.e.

$$F_\theta(i,j) = (K_\theta * I)(i,j) \qquad (1)$$

where (i,j) is the pixel location, and θ is the orientation angle. For each pixel location (i,j), we pick up the θ which maximizes $F_\theta(i,j)$, and the corresponding maximal value to obtain an orientation map $I_\theta(i,j)$ and a texture map $I_t(i,j)$,
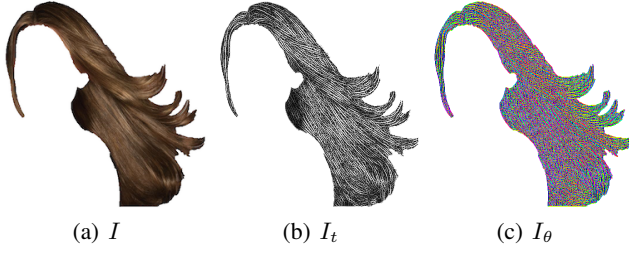
(a) I (b) I_t (c) I_θ

Figure 4: Visualization of texture map and orientation map. (a) Input image. (b)(c) Texture map and orientation map of (a) respectively. Here orientation map is colorized for visualization.

respectively, i.e.

$$I_\theta(i,j) = \arg\max_\theta\{F_\theta(i,j)\} \qquad (2)$$

$$I_t(i,j) = \max_\theta\{F_\theta(i,j)\} \qquad (3)$$

In our paper, we utilize 8 even-symmetric cosine Gabor kernels as the filter bank $\{K_\theta\}$, where the orientation angle θ being evenly sampled between $0°$ and $180°$. Specifically, the cosine Gabor kernel at angle θ is defined as:

$$K_\theta(u,v) = \exp\left(-\frac{1}{2}\left[\frac{\tilde{u}^2}{\sigma_u^2} + \frac{\tilde{v}^2}{\sigma_v^2}\right]\right) \cdot \cos\left(\frac{2\pi\tilde{u}}{\lambda}\right) \qquad (4)$$

where $\tilde{u} = u\cos\theta + v\sin\theta$ and $\tilde{v} = -u\sin\theta + v\cos\theta$. σ_u, σ_v, λ are hyper-parameters, and we simply set them 1.8, 2.4, 4 in all of our experiments respectively. We denote the operations stated above as $[I_\theta, I_t] \leftarrow \mathbf{g}(I)$. Figure 3 illustrates Gabor responses in 4 orientations, and Figure 4 visualizes an example of texture map and orientation map.

In practice, we found that I_t initially extracted from I_c contains most textures but some blurry artifacts in I_c also retains. To tackle this issue, we duplicate the operation $\mathbf{g}$, first applying it to I_c to produce I_t^0 and then re-applying it to I_t^0 to obtain the final I_t and I_θ, i.e.

$$[I_\theta, I_t] \leftarrow \mathbf{g}(I_t^0), \quad \text{where} \quad [I_\theta^0, I_t^0] \leftarrow \mathbf{g}(I_c) \qquad (5)$$

We demonstrate the effectiveness of using $[I_t, I_\theta]$ instead of $[I_t^0, I_\theta^0]$ in Figure 5.

Network Structure and Optimization Goal
$[I_t, I_\theta]$ obtained by the structure extraction layer now contains not only the structure but also the detailed information of hair. To take fully advantage of these high-level and low-level features, we apply U-Net architecture [Ronneberger *et al.*, 2015] which is composed of fully convolutional neural networks and skip connections, and produce $I_f = G_r(I_c)$. Also, the optimization goal is formulated as a combination of several types of losses, including pixel loss L_{pixel}, adversarial loss L_{adv}, style loss L_{style}, and feature matching loss L_{FM}.

Pixel-level reconstruction. The output produced by our regenerating network needs to respect the ground truth in pixel-level as the traditional methods do [Radford *et al.*, 2015; Isola *et al.*, 2017]. This is achieved by computing the L_1 loss between I_f and ground truth image I_{gt} as in [Wang *et*

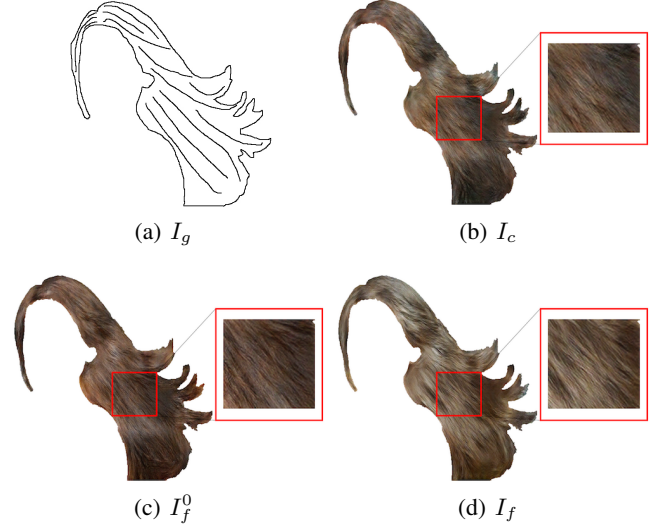(a) I_g (b) I_c

(c) I_f^0 (d) I_f

Figure 5: Comparison of the final results, by using structure information $[I_t^0, I_\theta^0]$ and $[I_t, I_\theta]$. (a) An input sketch image. (b) The coarse result generated by basic network G_b. (c)(d) The final results generated by using the extracted structure information $[I_t^0, I_\theta^0]$ and $[I_t, I_\theta]$, respectively. Compared with (c), (d) is in higher quality in terms of textures and gloss.

al., 2018a; Meng *et al.*, 2018; Wang *et al.*, 2019]. Similarly, to encourage sharp and varied hair generation, we also introduce adversarial loss L_{adv} as in [Goodfellow *et al.*, 2014]. These two losses are written as

$$L_{pixel} = \frac{1}{H \times W \times C}\|I_{gt} - G_r(I_c)\|_1 \qquad (6)$$

$$L_{adv} = -\sum_{I_c}\log D(G_r(I_c)) \qquad (7)$$

where H, W, C are the sizes of the image, and D is a discriminator of GAN.

Style reconstruction. One core issue of synthesizing plausible hair image is generating silky and glossy hair style. Training with L_{pixel} and L_{adv} forces the network paying too much attention to the per-pixel color transfer and limited high-frequency generation, instead of reconstructing realistic styles of hair. Therefore, we incorporate the associated constraint by measuring the input-output distance in a proper feature space, which is sensitive to the style change such as integral colors, textures and common patterns, while relatively robust to other variations. Specifically, we define L_{style} the following manner.

Denote ψ as a pre-trained network for feature extraction and $\psi^i(x)$ is the feature map at its i_{th} layer, whose shape is $C_i \times H_i \times W_i$. A Gram matrix of shape $C_i \times C_i$ as introduced in [Gatys *et al.*, 2016] is built with its entry defined as:

$$\mathbf{G}^i(x)_{j,k} = \frac{1}{C_iH_iW_i}\sum_{h=1}^{H_i}\sum_{w=1}^{W_i}\psi^i(x)_{j,h,w} \cdot \psi^i(x)_{k,h,w} \qquad (8)$$

Then the style reconstruction loss at the i_{th} layer is formulated as the squared Frobenius norm of the difference between

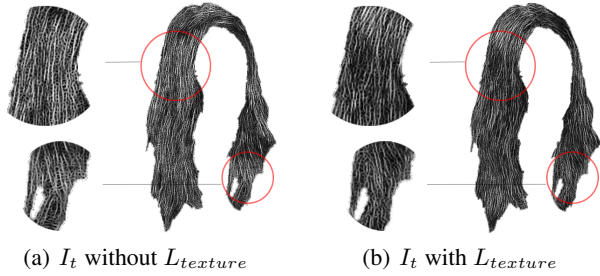(a) I_t without $L_{texture}$ (b) I_t with $L_{texture}$

Figure 6: Texture map extracted without vs. with $L_{texture}$. Compared with (a), clearer hairlines are generated in (b) with the guide of texture loss.

the Gram matrices of the output and target image:

$$L^i_{style}(I_c, I_{gt}) = \left\| \mathbf{G}^i(G_r(I_c)) - \mathbf{G}^i(I_{gt}) \right\|^2_F \qquad (9)$$

In practice, we apply a pre-trained VGG-16 network as ψ, and accumulate the style losses on its two layers (relu2_2, relu3_3) to formulate L_{style}. In addition, for efficiency, we reshape $\psi^i(x)$ into a matrix χ of shape $C_i \times H_i W_i$, so as to obtain $\mathbf{G}^i(x) = \chi\chi^T / C_i H_i W_i$.

Feature matching. As discussed in [Wang *et al.*, 2018b], high-resolution image synthesis poses a great challenge to the GAN discriminator design. To differentiate high-resolution real and synthesized images, the discriminator D needs to be deeper to have a large receptive field. Thus it is difficult for D to penalize the difference in detail. A featuring matching loss [Wang *et al.*, 2018b] defined on the feature layers extracted by D can well simplify this problem, i.e.

$$L_{FM}(G_r, D) = \mathbb{E}_{I_c} \sum_{i=1}^{T} \frac{1}{N_i} \left\| D^i(I_{gt}) - D^i(G_r(I_c)) \right\|_1$$
$$(10)$$

where D^i is the feature map in the i_{th} layer of D, N_i is the number of elements in each layer and T is the total number of layers.

Our full objective combines all of the losses above as

$$\min_{G_r} \left(w_1 L_{pixel} + w_2 L_{adv} + w_3 L_{style} + w_4 L_{FM} \right) \quad (11)$$

where $w_1 \sim w_4$ controls the importance of the 4 terms.

3.3 Training Strategy

The basic network and the re-generating network are trained separately at the very beginning. For the re-generating network, we first directly extract the structure information $[I_t, I_\theta]$ from the ground truth I_{gt}. When recognizable coarse images are available from the trained basic network, we replace the source of $[I_t, I_\theta]$ from ground truth image to the coarse image I_c. This strategy enables the trained re-generating network to see enough data to avoid easily overfitting. Finally, we connect the two networks and fine-tune them jointly with a reduced learning rate.
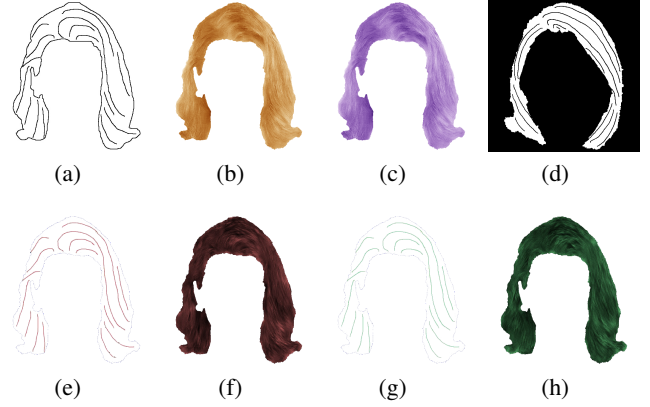

Figure 7: Colorized results by two solutions, (a) to (d) by postprocessing, i.e. tuning tones in Photoshop, and (e) to (h) by data-driven approach. (e) and (g) are the input sketch images with color strokes, (f) and (h) are their corresponding outputs.

4 Applications

4.1 Sketch2Hair

The input for the Sketch2Hair task is a binary image I_g with 2 channels. Its first channel is a mask $M \in \{0, 1\}$ where 0 and 1 fill the outside and inside of the contour of the hair, respectively. Its second channel $S \in \{0, 1\}$ encodes the strokes (0) indicating the dominant growth directions of strands. As aforementioned, in Phase 1, I_g is fed into a basic generating network G_b to produce a coarse result I_c, and then a re-generating network G_r is applied to produce the final result I_f.

For G_b, we choose a deep U-net [Ronneberger *et al.*, 2015] that convolves the input sketch into 1×1 feature maps and then deconvolutes it iteratively to generate I_c of the same size as I_g. The reason behind is that, we require the network has enough capability to imagine the whole hairlines, given the rough flow information in the sparse strokes. It should not trivially view the strokes as the only hairlines for the synthesized image, in which case, the network may only render some background color between them as in [Chen and Koltun, 2017]. Deconvolution from 1×1 feature maps makes each recovered pixel sees all input pixels, and deep structure makes the network has more parameters to learn from the data.

Additionally, during training G_b, we add a texture loss $L_{texture}$ into the existing L_{pixel} and L_{adv}, to enforce I_c containing as much as texture information as possible. Specifically, texture loss penalizes the texture difference between I_{gt} and I_c, and is written as

$$L_{texture} = \left\| I_t(I_{gt}) - I_t(I_c) \right\|_1 \qquad (12)$$

where $I_t(x)$ means the texture map extracted from x as defined similarly in Eq. (5). We illustrate two texture maps extracted with vs. without $L_{texture}$ after optimizing G_b in Figure 6. As seen, with this loss included, clearer hairlines are present in the texture map, facilitating the high-quality image generation in the successive network.

Figure 8: A glance of the hair dataset, including HD hair images (odd columns) and the corresponding sketches (even columns).

Then we attach G_r to the end of G_b to improve the final results. We list extensive comparisons in Section 5 to demonstrate the effectiveness of G_r, specifically in Figure 1 and 9.

I_f generated by G_r is usually drab and lacks of rich colors, while a color-specified synthesized hair image is often required. To this end, we provide two solutions for it. One is to tune the tone of I_f in image processing software like Photoshop, which preserves the gloss and silky style as in I_f. The other one is data-driven similar to [Sangkloy *et al.*, 2017], i.e. augment the training pairs into colorized version, so as to enable the networks to learn the color correspondence. Experimental results in Figure 7 show that the two solutions achieve pleasing outputs.

4.2 Hair Super-Resolution

Super-resolution for human portrait images is in higher demand due to the increasing use of digital camera apps, where hair super-resolution plays a key role. In this task, the input image I_g is a low-resolution image and I_f is in higher resolution. Considering visual performance, we choose SR-GAN [Ledig *et al.*, 2017] and ESRGAN [Wang *et al.*, 2018c] as our basic nets to produce I_c. In experiments, we found that if the input image is too small, the basic nets will fail in producing accurate results, obtaining blur artifacts. With our re-generating network introduced, its self-enhancing capability enables the generation of finer textures and fewer artifacts exist in I_f. Note that here, we also feed a bicubic upsampled version $\tilde{I}_g$ of input I_g to G_r, because $[I_t, I_\theta]$ contains only structure instead of color information.

We also list the comparison results in Figure 10 to demonstrate the effectiveness of our approach.

5 Experimental Results

5.1 Dataset Construction

We built our dataset in the following manner. First, we collected 640 high-resolution photographs with clear hair strands from internet, where the hair area is restricted up to 512×512. Then for each image, the hair region is manually segmented and the mask is saved. For Sketch2Hair, we distributed the 640 images to 5 paid users, who were requested to draw a set of lines in the hair area according to the growth directions of strands, generating 640 sketch-image pairs. For super-resolution, we only downsampled the 512^2 resolution images

pix2pix	CRN	pix2pixHD	pix2pixStyle	ours
7.5943	9.7551	7.4186	7.4069	**7.3827**

	4× (SRGAN as G_b)	8× (ESRGAN as G_b)
I_c	9.5084	7.7891
I_f (ours)	**7.5229**	**7.6702**

Table 1: NIQE scores of Sketch2Hair (top) and super-resolution results (bottom). Lower value represents better quality.

to obtain the input data. We randomly split the whole data into training and testing sets with a ratio of $4:1$. A glance of the dataset is illustrated in Figure 8.

5.2 Comparisons on Sketch2Hair

For the task of Sketch2Hair, we based our comparisons on pix2pix [Isola *et al.*, 2017] , pix2pixHD [Wang *et al.*, 2018b], CRN [Chen and Koltun, 2017] and pix2pixStyle. For pix2pix, we just used the official code while replacing U-net256 with U-net512. CRN is a single feedforward network without adversarial structure. We firstly trained it in 256^2 resolution and then fine-tuned it in 512^2 resolution. Pix2pixHD is an improved model of pix2pix for large scale image synthesis, so we also used its official code on 512^2 resolution. For Pix2pixStyle, we trained it on the structure of pix2pix with an additional style loss. For fair comparison, the 4 models were trained thoroughly with our dataset, until the convergence of their losses.

Figure 9 shows the qualitative comparisons of our approach against the above 4 methods. As seen, the basic pix2pix framework produced blurry artifacts. Pix2pixHD improved the appearance by better generating the textured structure, while noise still exists in the results. CRN produced the most visual-pleasing outputs, but fine details were missing. Pix2pixStyle has similar performances as pix2pixHD while it lacks of enough high frequency information. By contrast, our approach not only synthesized better structure but also generated finer details, producing a silky and glowing appearance.

We also applied a no-reference image quality score i.e. Naturalness Image Quality Evaluator (NIQE) to quantitatively measure the quality of the results. A smaller score indicates better perceptual quality. The evaluated scores demonstrate our approach outperforms the 4 models above as seen in Table 1. We further conducted a user study to perceptually evaluate the performances of the models stated above. Specifically, we randomly picked up 20 images from the total 128 testing images and paired them with the corresponding outcomes of the above 4 approaches, obtaining 80 pairs. Then we shuffled these pairs and asked volunteers to choose the one with finer strands in each pair. 65 volunteers were involved, 48 of whom are male, 94% are from 18-25 years old, and 6% are from 26-35. The results show that in the comparisons with CRN, pix2pix, pix2pixHD, and pix2pixStyle, our approach was voted by 98.23%, 94.38%, 86.62%, 69.62% of volunteers separately.

5.3 Comparisons on Hair Super-Resolution

For the task of Hair Super-Resolution, we compared our approach with two state-of-the-art methods, SRGAN [Ledig *et al.*, 2017] and ESRGAN [Wang *et al.*, 2018c], in 4x scale

(a) input (b) pix2pix (c) CRN (d) pix2pixHD (e) pix2pixStyle (f) ours (g) ground truth

Figure 9: Qualitative comparisons of our method against alternative approaches, for the task of Sketch2Hair.

and 8x scale separately. We list the results in Figure 10. As seen, SRGAN recovered limited image details in 4x results but the details are still fuzzy. ESRGAN produced clear enough but gummy hairlines in the 8x results. In contrast, our approach significantly reduced these artifacts and accurately reconstructed the finer details, producing plausible and visually pleasant results.

Similarly, NIQE was also applied here to quantitatively measure the quality of the results. We list the scores of results I_c by SRGAN and ESRGAN, as well as the enhanced results I_f by our method in Table. 1, which demonstrate the self-enhancing capability of our re-generating network.

6 Conclusion and Discussion

We have presented a hair image synthesis approach given an image with limited guidance, such as sketch or low-resolution image. We generalize the hair synthesis in two phases. In Phase 1, we apply an existing image-to-image translation network to generate a coarse result, which is recognizable but lacks of textured details. Then we apply a re-generating network with self-enhancing capability to the coarse result, and produce the final high-quality result. The self-enhancing capability is achieved by a proposed structure extraction layer, which extracts the texture and orientation map from a hair image. Experimental results demonstrated that our method outperforms state-of-the-art, in perceptual user study, qualitative and qualitative comparisons. We hope that this two-phase approach could be potentially applied to more hair-synthesis works.

(a) I_g (b) I_c (c) I_f (d) I_{gt}

Figure 10: Results by our approach and SRGAN / ESRGAN in the task of hair image super-resolution. (a) Low-resolution input. (b) SR results in 4x by SRGAN (top 2 rows) and in 8x by ESRGAN (bottom 2 rows). (c) Our enhanced results from (b). (d) Ground-truth. Enhanced hair are clearer and more visually pleasant (Better watching via zooming in).

Our results, while significantly outperforms the state-of-the-art in the realism, are still distinguishable from real photographs. Also, our method lacks the ability to control the shading and depth information of the generated hair, as no such information is available in the input sketch. Also, our method cannot handle complex hairstyles such as braids. Techniques like combining multi-modal features together as in [Wang *et al.*, 2014; Ren *et al.*, 2016] could be potential solutions. We will leave all of the goals as our future work.

References

[Cao *et al.*, 2017] Qingxing Cao, Liang Lin, Yukai Shi, Xiaodan Liang, and Guanbin Li. Attention-aware face hallucination via deep reinforcement learning. pages 1656–1664, 2017.

[Chai *et al.*, 2012] Menglei Chai, Lvdi Wang, Yanlin Weng, Yizhou Yu, Baining Guo, and Kun Zhou. Single-view hair modeling for portrait manipulation. *ACM Transactions on Graphics (TOG)*, 31(4):116, 2012.

[Chai *et al.*, 2013] Menglei Chai, Lvdi Wang, Yanlin Weng, Xiaogang Jin, and Kun Zhou. Dynamic hair manipulation in images and videos. *ACM Transactions on Graphics (TOG)*, 32(4):75, 2013.

[Chai *et al.*, 2016] Menglei Chai, Tianjia Shao, Hongzhi Wu, Yanlin Weng, and Kun Zhou. Autohair: fully automatic hair modeling from a single image. *ACM Transactions on Graphics*, 35(4), 2016.

[Chen and Koltun, 2017] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *IEEE International Conference on Computer Vision (ICCV)*, page 3, 2017.

[Chen and Zhu, 2006] Hong Chen and Song-Chun Zhu. A generative sketch model for human hair analysis and synthesis. *IEEE transactions on pattern analysis and machine intelligence*, 28(7):1025–1040, 2006.

[Dong *et al.*, 2014] Chao Dong, Change Loy Chen, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. 8692:184–199, 2014.

[Dong *et al.*, 2016] Chao Dong, Change Loy Chen, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 38(2):295–307, 2016.

[Fu *et al.*, 2007] Hongbo Fu, Yichen Wei, Chiew-Lan Tai, and Long Quan. Sketching hairstyles. In *Proceedings of the 4th Eurographics workshop on Sketch-based interfaces and modeling*, pages 31–36. ACM, 2007.

[Gatys *et al.*, 2016] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016.

[Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[Hu *et al.*, 2015] Liwen Hu, Chongyang Ma, Linjie Luo, and Hao Li. Single-view hair modeling using a hairstyle database. *ACM Transactions on Graphics (Proceedings SIGGRAPH 2015)*, 34(4), July 2015.

[Huang *et al.*, 2017] Huaibo Huang, Ran He, Zhenan Sun, and Tieniu Tan. Wavelet-srnet: A wavelet-based cnn for multi-scale face super resolution. In *IEEE International Conference on Computer Vision*, pages 1698–1706, 2017.

[Isola *et al.*, 2017] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017.

[Jakob *et al.*, 2009] Wenzel Jakob, Jonathan T Moon, and Steve Marschner. Capturing hair assemblies fiber by fiber. *ACM Transactions on Graphics (TOG)*, 28(5):164, 2009.

[Karras *et al.*, 2017] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.

[Kim *et al.*, 2016] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Computer Vision and Pattern Recognition*, pages 1646–1654, 2016.

[Lai *et al.*, 2009] Yu-Kun Lai, Shi-Min Hu, and Ralph R Martin. Automatic and topology-preserving gradient mesh generation for image vectorization. In *ACM Transactions on Graphics (TOG)*, volume 28, page 85. ACM, 2009.

[Lassner *et al.*, 2017] Christoph Lassner, Gerard Pons-Moll, and Peter V. Gehler. A generative model for people in clothing. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.

[Ledig *et al.*, 2017] Christian Ledig, Zehan Wang, Wenzhe Shi, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, and Alykhan Tejani. Photo-realistic single image super-resolution using a generative adversarial network. In *Computer Vision and Pattern Recognition*, pages 105–114, 2017.

[Levinshtein *et al.*, 2018] Alex Levinshtein, Cheng Chang, Edmund Phung, Irina Kezele, Wenzhangzhi Guo, and Parham Aarabi. Real-time deep hair matting on mobile devices. 2018.

[Li *et al.*, 2018] Xiaoming Li, Ming Liu, Yuting Ye, Wangmeng Zuo, Liang Lin, and Ruigang Yang. Learning warped guidance for blind face restoration. 2018.

[Lim *et al.*, 2017] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1132–1140, 2017.

[Liu *et al.*, 2016] Sifei Liu, Jinshan Pan, and Ming Hsuan Yang. *Learning Recursive Filters for Low-Level Vision via a Hybrid Neural Network*. Springer International Publishing, 2016.

[Mao *et al.*, 2004] Xiaoyang Mao, Hiroki Kato, Atsumi Imamiya, and Ken Anjyo. Sketch interface based expressive hairstyle modelling and rendering. In *Computer Graphics International, 2004. Proceedings*, pages 608–611. IEEE, 2004.

[Meng *et al.*, 2018] Xiandong Meng, Xuan Deng, Shuyuan Zhu, Shuaicheng Liu, Chuan Wang, Chen Chen, and Bing Zeng. Mganet: A robust model for quality enhancement of compressed video. *arXiv preprint arXiv:1811.09150*, 2018.

[Mirza and Osindero, 2014] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.

[Pan *et al.*, 2018] Jinshan Pan, Sifei Liu, Deqing Sun, Jiawei Zhang, Yang Liu, Jimmy Ren, Zechao Li, Jinhui Tang, Huchuan Lu, and Yu Wing Tai. Learning dual convolutional neural networks for low-level vision. 2018.

[Portenier *et al.*, 2018] Tiziano Portenier, Qiyang Hu, Attila Szabó, Siavash Arjomand Bigdeli, Paolo Favaro, and Matthias Zwicker. Faceshop: Deep sketch-based face image editing. *ACM Trans. Graph.*, 37(4):99:1–99:13, July 2018.

[Radford *et al.*, 2015] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

[Ren *et al.*, 2016] Jimmy Ren, Yongtao Hu, Yu-Wing Tai, Chuan Wang, Li Xu, Wenxiu Sun, and Qiong Yan. Look, listen and learnła multimodal lstm for speaker identification. In *30th AAAI Conference on Artificial Intelligence*, 2016.

[Ronneberger *et al.*, 2015] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *medical image computing and computer assisted intervention*, pages 234–241, 2015.

[Sangkloy *et al.*, 2017] Patsorn Sangkloy, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. Scribbler: Controlling deep image synthesis with sketch and color. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6836–6845, 2017.

[Shelhamer *et al.*, 2017] E. Shelhamer, J. Long, and T. Darrell. Fully convolutional networks for semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):640–651, April 2017.

[Wang *et al.*, 2014] Chuan Wang, Yanwen Guo, Jie Zhu, Linbo Wang, and Wenping Wang. Video object co-segmentation via subspace clustering and quadratic pseudo-boolean optimization in an mrf framework. *IEEE Transactions on Multimedia*, 16(4):903–916, 2014.

[Wang *et al.*, 2017] Chuan Wang, Jie Zhu, Yanwen Guo, and Wenping Wang. Video vectorization via tetrahedral remeshing. *IEEE Transactions on Image Processing*, 26(4):1833–1844, 2017.

[Wang *et al.*, 2018a] Chuan Wang, Haibin Huang, Xiaoguang Han, and Jue Wang. Video inpainting by jointly learning temporal structure and spatial details. *arXiv preprint arXiv:1806.08482*, 2018.

[Wang *et al.*, 2018b] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, page 5, 2018.

[Wang *et al.*, 2018c] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Chen Change Loy, Yu Qiao, and Xiaoou Tang. Esrgan: Enhanced super-resolution generative adversarial networks. 2018.

[Wang *et al.*, 2019] Yang Wang, Haibin Huang, Chuan Wang, Tong He, Jue Wang, and Minh Hoai. Gif2video: Color dequantization and temporal interpolation of gif images. *arXiv preprint arXiv:1901.02840*, 2019.

[Wei *et al.*, 2005] Yichen Wei, Eyal Ofek, Long Quan, and Heung-Yeung Shum. Modeling hair from multiple views. In *ACM Transactions on Graphics (ToG)*, volume 24, pages 816–820. ACM, 2005.

[Wei *et al.*, 2018] Lingyu Wei, Liwen Hu, Vladimir Kim, Ersin Yumer, and Hao Li. Real-time hair rendering using sequential adversarial networks. In *The European Conference on Computer Vision (ECCV)*, September 2018.

[Weng *et al.*, 2013] Yanlin Weng, Lvdi Wang, Xiao Li, Menglei Chai, and Kun Zhou. Hair interpolation for portrait morphing. In *Computer Graphics Forum*, volume 32, pages 79–84. Wiley Online Library, 2013.

[Xu *et al.*, 2015] Li Xu, Jimmy S. J Ren, Qiong Yan, Renjie Liao, and Jiaya Jia. Deep edge-aware filters. In *International Conference on International Conference on Machine Learning*, pages 1669–1678, 2015.

[Yu and Porikli, 2016] Xin Yu and Fatih Porikli. Ultra-resolving face images by discriminative generative networks. 2016.

[Yu and Porikli, 2017] Xin Yu and Fatih Porikli. Face hallucination with tiny unaligned images by transformative discriminative neural networks, 2017.

[Zhang *et al.*, 2018] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. 2018.

[Zheng *et al.*, 2015] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. S. Torr. Conditional random fields as recurrent neural networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1529–1537, Dec 2015.

[Zhou *et al.*, 2018] Yi Zhou, Liwen Hu, Jun Xing, Weikai Chen, Han-Wei Kung, Xin Tong, and Hao Li. Single-view hair reconstruction using convolutional neural networks. *arXiv preprint arXiv:1806.07467*, 2018.

[Zhu *et al.*, 2016] Shizhan Zhu, Sifei Liu, Chen Change Loy, and Xiaoou Tang. Deep cascaded bi-network for face hallucination. pages 614–630, 2016.