

HairFIT: Pose-invariant Hairstyle Transfer via Flow-based Hair Alignment and Semantic-region-aware Inpainting

Chaeyeon Chung*^{1,2}

cy_chung@kaist.ac.kr

Taewoo Kim*^{1,2}

specia1ktu@kaist.ac.kr

Hyelin Nam*³

nhl0208@cau.ac.kr

Seunghwan Choi¹

shadow2496@kaist.ac.kr

Gyojung Gu¹

gyojung.gu@kaist.ac.kr

Sunghyun Park^{1,2}

psh01087@kaist.ac.kr

Jaegul Choo¹

jchoo@kaist.ac.kr

¹ KAIST AI

South Korea

² Kakao Enterprise

South Korea

³ Chung-Ang University

South Korea

Abstract

Hairstyle transfer is the task of modifying a source hairstyle to a target one. Although recent hairstyle transfer models can reflect the delicate features of hairstyles, they still have two major limitations. First, the existing methods fail to transfer hairstyles when a source and a target image have different poses (*e.g.*, viewing direction or face size), which is prevalent in the real world. Also, the previous models generate unrealistic images when there is a non-trivial amount of regions in the source image occluded by its original hair. When modifying long hair to short hair, shoulders or backgrounds occluded by the long hair need to be inpainted. To address these issues, we propose a novel framework for pose-invariant hairstyle transfer, HairFIT. Our model consists of two stages: 1) flow-based hair alignment and 2) hair synthesis. In the hair alignment stage, we leverage a keypoint-based optical flow estimator to align a target hairstyle with a source pose. Then, we generate a final hairstyle-transferred image in the hair synthesis stage based on Semantic-region-aware Inpainting Mask (SIM) estimator. Our SIM estimator divides the occluded regions in the source image into different semantic regions to reflect their distinct features during the inpainting. To demonstrate the effectiveness of our model, we conduct quantitative and qualitative evaluations using multi-view datasets, K-hairstyle and VoxCeleb. The results indicate that HairFIT achieves a state-of-the-art performance by successfully transferring hairstyles between images of different poses, which have never been achieved before.

* indicates equal contributions.

©2021. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

1 Introduction

Recently, hairstyle has been considered as a way to express one’s own identity. Responding to the increasing demands, various approaches have tackled virtual hairstyle transfer, a task of modifying one’s hairstyle to a target one, based on generative adversarial networks [8]. The existing methods [25, 28] transfer the target hairstyle to a source image while successfully preserving delicate features of the target hairstyle. However, despite their convincing performances, hairstyle transfer still poses several challenges.

First, the existing models are hardly applicable to a source image that has a different pose (*e.g.*, viewing direction or face size) from a target image, which is prevalent in the real world. The first two rows of Fig. 3 present examples of the pose difference between the source and the target image. Since the same hairstyle appears different when its pose changes, the target hairstyle should have a similar pose with the source image. However, the previous approaches fail to align the target hair with the source pose when their poses are significantly different.

Additionally, the existing methods poorly generate regions originally occluded in the source image. When the source image contains the regions that are occluded by its original hair, the occlusions need to be newly generated during the transfer. For instance, the first row of Fig. 4 shows that the occluded shoulders and backgrounds should be properly inpainted when transferring the shorter hair. Moreover, the occlusions become even larger when the source and the target image have a significant pose difference. Although the existing models leverage the external inpainting networks, it is still challenging for a model to fill in the occlusions containing multiple semantic regions.

To address these issues, we propose a novel framework, Hairstyle transfer via Flow-based hair alignment and semantic-region-aware Inpainting (HairFIT), which successfully performs hairstyle transfer regardless of a pose difference between a source and a target image. HairFIT consists of two parts: 1) flow-based hair alignment and 2) hair synthesis. We first align the target hair with the source pose based on the estimated optical flow between the source and the target. Then, we apply the aligned target hair to the source image to synthesize the final output. The hair synthesis module inpaints in the occluded regions of the source image and refines the aligned target hair leveraging ALIGNment-Aware Segment (ALIAS) generator [9]. We newly present a Semantic-region-aware Inpainting Mask (SIM) estimator to support realistic occlusion inpainting. By predicting separate semantic regions, the estimator helps our generator to reflect distinctive features of diverse semantic regions (*e.g.*, backgrounds, clothes, and forehead) during inpainting.

We conduct quantitative and qualitative evaluations using K-hairstyle and VoxCeleb which contain images of different poses. The results indicate that our model achieves a state-of-the-art performance compared to the existing methods. Our model outperforms other models especially when the source and the target image have a significantly different pose.

Our contributions are summarized as follows:

- This is the first work that proposes a pose-invariant hairstyle transfer framework by adopting the optical flow estimation for hair alignment.
- We present a Semantic-region-aware Inpainting Mask (SIM) estimator which supports high-quality occlusion inpainting by allowing hair synthesis module to reflect distinctive features of each segmented inpainting region.
- HairFIT achieves the state-of-the-art performance in both quantitative and qualitative evaluations with two multi-view datasets, K-hairstyle and VoxCeleb.

2 Related Work

2.1 Conditional Image Generation

Conditional image generation is the task of synthesizing images based on the given conditions such as category labels [1, 19], text [24, 36], and images [12, 21]. Recent studies on image-conditioned generation have proposed various approaches to modify specific features (e.g., face shape, hairstyle, etc.) on facial images [13, 15, 23, 32]. In particular, MaskGAN [15] enables facial image manipulation based on a user-edited semantic mask as a conditional input. Also, SC-FEGAN [13], FaceShop [23], and Deep Plastic Surgery [32] allow modification of facial attributes based on human-drawn sketches. However, the existing methods require a non-trivial amount of user interaction to obtain the desired conditional inputs to change specific target attributes. In this paper, we propose a framework that transfers a hairstyle without any user-modified conditional input or user interaction.

2.2 Hairstyle Transfer

Hairstyle transfer aims to modify a hairstyle of a source image to a target one while preserving other features of the source. Along with the researches on facial attribute modification [13, 23, 32], hairstyle transfer has also been actively tackled in the recent few years [25, 28]. MichiGAN [28] successfully transfers a hairstyle reflecting its delicate features using disentangled hair attributes and conditional synthesis modules of each attribute. In addition, LOHO [25] generates realistic hairstyle-transferred images by decomposing hairstyle features and optimizing latent space to inpaint missing hair structure details. However, these works only tackle the cases where poses of a source and a target image are similar. Therefore, they fail to generate realistic images when the source and the target have significantly different viewing directions or face sizes. Unlike the existing approaches, our model successfully transfers a hairstyle even when the source and the target image have a significant pose difference.

2.3 Optical Flow Estimation

Optical flow represents apparent pixel-level motion patterns between two images using a vector field. There are several approaches based on neural networks that directly predict an optical flow map between two consecutive frames of a video [9, 11, 14] and warp the raw pixels of a frame to generate the unseen frame [5, 16, 17, 21]. The optical flow estimation is also leveraged to transfer motions from another video by learning the motion difference between a source and a target frame [26, 27]. To achieve pose-invariant hairstyle transfer, we utilize the optical flow to deform the target hairstyle according to the pose of the source image. To the best of our knowledge, this is the first work that applies optical flow estimation to hairstyle transfer.

3 Method

3.1 Overview

We employ a two-stage approach with two separate modules: 1) flow-based hair alignment module and 2) hair synthesis module. When given a source image and a target image con-

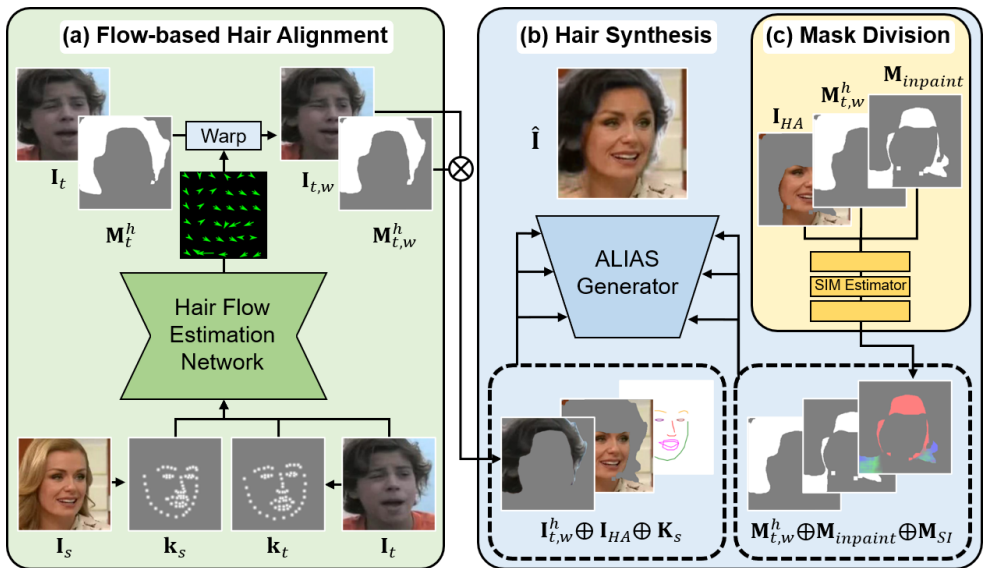


Figure 1: Overall architecture of HairFIT. Our model consists of two modules: (a) flow-based hair alignment module and (b) hair synthesis module. (a) Hair alignment module aligns the hair from a target image I_t with the pose of a source image I_s based on facial keypoints k_s , k_t extracted from each image. In (b) hair synthesis stage, ALIAS generator synthesizes the final image \hat{I} based on (c) mask division with SIM estimator.

taining the target hairstyle, the flow-based hair alignment module first warps the target image and its hair segmentation mask to align the target hairstyle with the source face. Here, we utilize the optical flow predicted by our hair flow estimation network. Then, the hair synthesis module generates the final image based on the aligned target hair and the source image via semantic-region-aware inpainting and hair refinement. Fig. 1 describes the overall architecture of HairFIT.

We train our model with multi-view datasets, which contain diverse views for each individual. In particular, we sample two images (*i.e.*, a source and a target image), which have different views with the same identity and hairstyle. Then, our model is required to reconstruct the source image by transferring the hairstyle of the target image to the source. Here, the source image works as the ground truth the model aims to generate. In this manner, the flow-based hair alignment module learns to align the target hair with the source pose according to the pose difference, and the hair synthesis module learns to reconstruct the original source image based on the aligned target hairstyle.

Ideally, when given a source and a target image with different poses, identities, and hairstyles, the ground truth image with the source pose and identity and the target hairstyle can provide HairFIT with direct supervision. In this case, the ground truth image can guide the model to align and apply the target hairstyle to the source face. However, a hairstyle dataset containing such images does not exist since it is expensive and time-consuming to collect them. Instead, we utilize the multi-view datasets for training our model. Although our model uses only image pairs of the same identity and hairstyle during the training, our model successfully transfers the hairstyle of the target image that has a different identity and hairstyle from the source.

3.2 Flow-based Hair Alignment

The flow-based hair alignment module aims to align the pose (*i.e.*, view and scale) of a target image with a source image.

As described in Fig. 1 (a), our hair flow estimation network estimates a dense optical flow map that represents a pose difference between the source image \mathbf{I}_s and the target image \mathbf{I}_t with the source keypoints \mathbf{k}_s , target keypoints \mathbf{k}_t , and \mathbf{I}_t given as inputs ($\mathbf{I} \in \mathbb{R}^{3 \times H \times W}$ and $\mathbf{k} \in \mathbb{R}^{N_k \times 2}$, where H , W , and N_k denote the image height, width, and the number of keypoints, respectively). Then, we warp \mathbf{I}_t as well as its hair segmentation mask \mathbf{M}_t^h according to the estimated optical flow. As a result, we gain the warped target hairstyle in $\mathbf{I}_{t,w}$ and its hair mask $\mathbf{M}_{t,w}^h$ aligned with the source face, where h indicates hair and w means the image or mask is warped.

Hair flow estimation network. Hair flow estimation network predicts the dense optical flow with the given \mathbf{k}_s , \mathbf{k}_t , and \mathbf{I}_t . To obtain the coarse pose difference, the network first converts \mathbf{k}_s and \mathbf{k}_t into a Gaussian keypoint heatmap $\mathbf{H} \in \mathbb{R}^{N_k \times H \times W}$, respectively. Then, we obtain the keypoint heatmap difference $\hat{\mathbf{H}}$, calculated as $\mathbf{H}_t - \mathbf{H}_s$. With $\hat{\mathbf{H}}$ and \mathbf{I}_t deformed by the keypoint difference $\hat{\mathbf{k}}$, which is $\mathbf{k}_t - \mathbf{k}_s$, the network predicts the N_k -channel mask $\mathbf{M} \in \mathbb{R}^{N_k \times H \times W}$ and \mathcal{F}_{ref} . \mathbf{M} contains the estimated local regions to apply each channel of $\hat{\mathbf{k}}$ to. Also, \mathcal{F}_{ref} indicates a refinement flow map that reflects additional detailed optical flow. Finally, we obtain the dense optical flow $\mathcal{F} \in \mathbb{R}^{2 \times H \times W}$ for the entire image by adding the coarse optical flow and refinement flow map as $\mathcal{F} = \sum_{i=1}^{N_k} \rho(\hat{\mathbf{k}}^i) \otimes \mathbf{M}^i + \mathcal{F}_{ref}$. Here, $\rho(\cdot)$ repeats the input tensor by $H \times W$ times and \otimes denotes element-wise multiplication.

More details of the network architectures are described in supplementary materials. For training, we use the reconstruction loss \mathcal{L}_{rec} , which is calculated as $\mathbb{E}[\|\mathbf{I}_{t,w} - \mathbf{I}_s\|_1]$. For the implementation, we referred to the existing keypoint-based optical flow estimation networks [26, 27].

Facial keypoints. Our hair flow estimation network utilizes keypoints to predict the flow map between the source and the target image. Here, the predicted optical flow needs to reflect the pixel-wise pose difference effectively, thus the keypoints should represent the overall view and scale of an image [27, 30]. In this regard, we utilize the facial keypoints, which consistently represent a person’s view (*e.g.*, head pose) or scale (*e.g.*, face size), regardless of his/her identity. Facial keypoints are simple but robust pose representations that guide the network to capture the key differences in poses. Since the datasets do not include ground truth facial keypoints, we extract 68 facial keypoints from each image using the pre-trained keypoint detector [2].

3.3 Hair Synthesis

The goal of the hair synthesis module is to transfer the target hairstyle to the source image using the aligned target hairstyle obtained from the previous stage. To achieve this, we utilize ALIAS generator [3] with a newly-proposed hair-agnostic image and a Semantic-region-aware Inpainting Mask (SIM) estimator.

Hair-agnostic image. At test time, our model needs to transfer hairstyle between the images of different identities and hairstyles, unlike the training phase. Therefore, the difference between the target hairstyle and the source hairstyle during the test is much larger than during the training. To minimize the difference, we obtain the hair-agnostic image \mathbf{I}_{HA} with the hair region of the source image \mathbf{I}_s removed as presented in Fig. 2. To be specific, we exclude the region of the source hair mask \mathbf{M}_s^h and the forehead mask \mathbf{M}_s^{fh} from \mathbf{I}_s . The

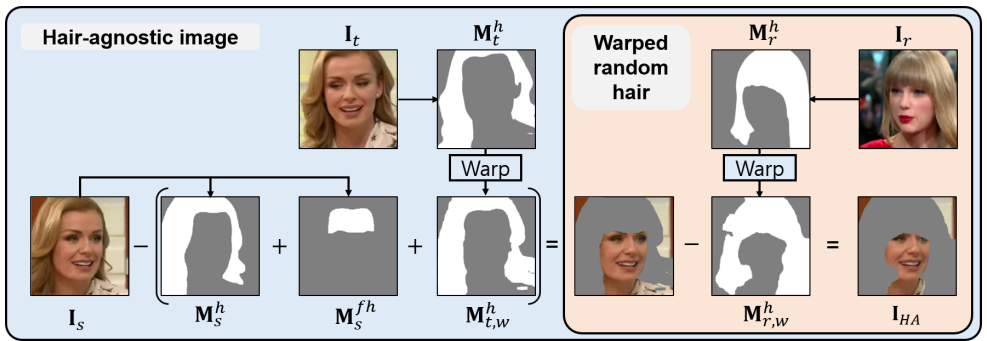


Figure 2: Process to obtain the hair-agnostic image \mathbf{I}_{HA} for hair synthesis. First, we remove the source hair region \mathbf{M}_s^h , the forehead region \mathbf{M}_s^{fh} , and the warped target hair region $\mathbf{M}_{t,w}^h$ from \mathbf{I}_s . We additionally remove a random warped hair region $\mathbf{M}_{r,w}^h$ during training to obtain \mathbf{I}_{HA} . $\mathbf{M}_{t,w}^h$ and $\mathbf{M}_{r,w}^h$ are warped by our flow-based hair alignment module.

warped target hair mask $\mathbf{M}_{t,w}^h$ is also subtracted from \mathbf{I}_s . During the training, $\mathbf{M}_{t,w}^h$ and \mathbf{M}_s^h have only a slight difference since we use image pairs of the same hairstyle. Therefore, we sample a random image \mathbf{I}_r from the training set and warp its hair mask \mathbf{M}_r^h to fit the source pose, resulting in $\mathbf{M}_{r,w}^h$. Then, the region of $\mathbf{M}_{r,w}^h$ is removed from the hair-agnostic image. This successfully mimics the inference situations by providing the model with the chance to fill in larger occluded regions. In this regard, the proposed \mathbf{I}_{HA} enhances our hair synthesis generator’s generalization ability at test time.

SIM estimator. As previously mentioned, the hairstyle transfer model is required to inpaint the occlusions in the source image and refine the aligned target hair. To address the challenges, we propose a SIM estimator, which separates the occluded regions into different semantic areas. With \mathbf{I}_{HA} , $\mathbf{M}_{t,w}^h$, and $\mathbf{M}_{inpaint}$ as inputs, SIM estimator predicts a semantic-region-aware inpainting mask $\mathbf{M}_{SI} \in \{0, 1\}^{4 \times H \times W}$ by dividing $\mathbf{M}_{inpaint} \in \{0, 1\}^{1 \times H \times W}$ into face, clothes, background, and unknown. Here, $\mathbf{M}_{inpaint}$ is calculated by subtracting $\mathbf{M}_{t,w}^h$ from the occluded region of \mathbf{I}_{HA} . The fifth column in Fig. 4 presents examples of \mathbf{M}_{SI} . The separated mask effectively guides our ALIAS generator to inpaint each semantic region by reflecting its distinctive features, leading to high-quality image generation.

ALIAS generator. As described in Fig. 1 (b), the inputs $\mathbf{I}_{t,w}^h \oplus \mathbf{I}_{HA} \oplus \mathbf{K}_s$ and $\mathbf{M}_{t,w}^h \oplus \mathbf{M}_{inpaint} \oplus \mathbf{M}_{SI}$ are injected into each layer of ALIAS generator, where $\mathbf{K}_s \in \mathbb{R}^{3 \times H \times W}$ is an RGB-rendered source keypoint image. In each layer, the features are normalized separately based on the inpainting mask $\mathbf{M}_{inpaint}$ in a similar manner to general region normalization [4, 35]. The normalized features are then modulated by β and γ , which are predicted based on $\mathbf{M}_{t,w}^h \oplus \mathbf{M}_{inpaint} \oplus \mathbf{M}_{SI}$. As a semantic guidance map of inpainting, \mathbf{M}_{SI} allows the modulation parameters β and γ to reflect distinctive semantic features of each divided region, increasing the overall quality of occlusion inpainting of the output. The hair synthesis module generates the final output after a series of ALIAS residual blocks with up-sampling layers. Additional details of ALIAS generator are provided in our supplementary material.

Losses. We train the ALIAS generator with the conditional adversarial loss \mathcal{L}_{cGAN} , the feature matching loss \mathcal{L}_{FM} , the perceptual loss $\mathcal{L}_{percept}$, and the hairstyle loss \mathcal{L}_{style} , referring to SPADE [22], pix2pixHD [29], and LOHO [23]. Since the training pairs have the same identity and hairstyle, the losses can be computed between a synthesized image and the ground

truth (source) image, working as direct supervision. We use the Hinge loss as the adversarial loss [6]. The \mathcal{L}_{style} is computed between Gram matrix [6] feature maps of the generated hair image and the source hair image to support detailed hairstyle refinement. The generated hair image is extracted from the final output using the pre-trained hair segmentation model [4]. Additionally, we train SIM estimator with SIM loss \mathcal{L}_{SIM} . \mathcal{L}_{SIM} is calculated using the binary cross-entropy loss between the estimated \mathbf{M}_{SI} and the ground truth segmentation masks obtained by the pre-trained face parsing model [3]. The total loss of the hair synthesis module is as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{cGAN} + \lambda_{FM}\mathcal{L}_{FM} + \lambda_{percept}\mathcal{L}_{percept} + \lambda_{style}\mathcal{L}_{style} + \lambda_{SIM}\mathcal{L}_{SIM}, \quad (1)$$

where we set both λ_{FM} and $\lambda_{percept}$ as 10, λ_{style} as 50, and λ_{SIM} as 100. More details of the losses are described in the supplementary material.

4 Experiments

To demonstrate the effectiveness of HairFIT, we conduct quantitative and qualitative evaluations compared to baseline models, using two multi-view datasets. We also conduct an ablation study to present the effect of key components in our model.

4.1 Experimental Setup

Dataset. As mentioned in Section 3.1, we leverage multi-view datasets for the experiments. First, we utilize K-hairstyle [4] which contains 500,000 high-resolution hairstyle images. The dataset consists of multi-view images including more than 6,400 identities. The viewing angle of images ranges from 0 to 360 degrees. Also, each image in the dataset has a hair segmentation mask as well as various hairstyle attributes. Since our goal is to transfer a hairstyle, we excluded images whose hairstyle is significantly occluded. Also, we removed images whose face is extremely rotated, such as images of a person facing the side or the back, since we cannot extract facial keypoints from them. Accordingly, our training set consists of randomly sampled 60,584 images with 5,407 identities, and the test set includes 6,717 images with 611 identities. While the maximum image resolution is $4,032 \times 3,024$, we resized the images into 256×256 and 512×512 in our experiments.

Additionally, we use VoxCeleb [8] which consists of more than 100,000 utterance videos with 1,251 identities. For the training data, we randomly sampled 16,847 videos, each of which contains 30 frames on average. The frames of a video can be considered as multi-view images of a single identity. The test set includes randomly sampled 2,209 videos. For the experiments, we resized the images into 256×256 .

For both datasets, the training set consists of image pairs that have the same identity and hairstyle, while the test set contains image pairs that have different identities and hairstyles. We utilize one image of a pair as a source image and the other as a target image.

Evaluation metrics. As a quantitative evaluation, we use the fréchet inception distance (FID) score [9]. The FID score measures how similar the distributions of the synthesized images and the real images are. The lower FID score indicates a higher similarity between the images.

Dataset	K-hairstyle		VoxCeleb
Resolution	256×256	512×512	256×256
MichiGAN	30.52	36.12	81.14
LOHO	62.44	71.79	66.97
HairFIT (Ours)	18.53	19.01	17.66

Table 1: Quantitative comparison with the baselines using K-hairstyle and VoxCeleb. We measure the FID scores.

4.2 Comparison to Baselines

Quantitative evaluations. First, we compare the FID scores between our model and the baseline models, MichiGAN [28] and LOHO [25]. We train both models with the same datasets, K-hairstyle and VoxCeleb, based on their official implementation codes. We utilize a gated convolution network [54] for the inpainting modules in MichiGAN and LOHO, as described in LOHO paper. As presented in Table 1, our model achieves the lowest FID score with a large margin compared to the baseline models.

Pose difference level	Medium	Difficult	Extremely difficult
MichiGAN	31.58	34.41	40.41
LOHO	57.80	62.34	95.89
HairFIT (Ours)	21.95	21.76	24.60

Table 2: Quantitative comparison with the baselines with three different levels of pose differences using K-hairstyle. We measure the FID scores.

For further analysis, we also conduct quantitative comparisons with three different levels of pose difference. We split our test pairs of K-hairstyle into three categories of 2,000 images, ‘Medium’, ‘Difficult’, and ‘Extremely difficult’. As in LOHO [25], we use 68 facial keypoints extracted by the pre-trained keypoint detector [9] to measure the pose distance (PD) between the source and the target keypoints as follows: $PD = \frac{1}{68} \sum_{i=1}^{68} \|\mathbf{k}_s - \mathbf{k}_t\|_2$. The result in Table 2 describes that HairFIT outperforms the other baselines with a larger margin as the degree of pose difference increases from ‘Medium’ to ‘Extremely difficult’. MichiGAN and LOHO show poor hairstyle transfer performance on images with significant pose differences, as they stated in their paper.

Additionally, to demonstrate that HairFIT successfully aligns the target hair with the source pose, we conduct a quantitative evaluation on hairstyle reconstruction. We compare

Dataset	K-hairstyle		VoxCeleb	
Metric	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow
MichiGAN	0.7210	0.2432	0.7021	0.2282
LOHO	0.7852	0.1452	0.6389	0.2419
HairFIT(Ours)	0.8041	0.0841	0.7147	0.1262

Table 3: Quantitative comparison with baselines on hairstyle reconstruction. We measure SSIM and LPIPS using K-hairstyle and VoxCeleb. \uparrow indicates the larger the better and \downarrow means the smaller the better.




Figure 3: Qualitative comparison with the baselines. (a) indicates the source images, (b) the target images, (c) the results of MichiGAN, (d) LOHO, and (e) ours. Due to the privacy issue, we blur the faces of the images from the K-hairstyle dataset.

our model to MichiGAN [28] and LOHO [25]. Given a source image and a target image which have the same identity and hairstyle but different poses, a model reconstructs the source image by applying the target hairstyle extracted from the target image to the source image. We measure the structural similarity (SSIM) [61] and learned perceptual image patch similarity (LPIPS) [68] using the same 3,000 image pairs. The higher the SSIM and the lower the LPIPS, the better. As described in Table 3, our model achieves superior performance over the baseline models.

Qualitative evaluations. As presented in Fig. 3, qualitative results also present the superiority of HairFIT for both of the two datasets. While the baseline models generate unrealistic images, our model robustly transfers target hairstyles to the source images even when they

Method	ALIAS only	+ Flow-based hair alignment	+ $M_{r,w}^h$ subtraction	+ SIM estimator with SIM loss	+ Hairstyle loss (Full)
FID (256)	25.58	22.89	19.40	19.16	18.53

Table 4: Ablation study on K-hairstyle dataset. Starting from ALIAS generator, we measure the FID scores by gradually adding each module and loss.



Figure 4: Qualitative evaluation on the SIM estimator. The first and the second row are examples of K-hairstyle and VoxCeleb, respectively. The red, green, and blue regions of M_{SI} indicate face, clothes, and background, respectively. Due to the privacy issue, we blur the faces of K-hairstyle images.

have significantly different poses or the source image has occluded regions.

4.3 Ablation Study

Quantitative evaluations. We conduct an ablation study to validate the effectiveness of each component in our model. Starting from ALIAS generator, we gradually add flow-based hair alignment module, $M_{r,w}^h$ subtraction in a hair-agnostic image I_{HA} , SIM estimator with SIM loss, and hairstyle loss, which is our full model. $M_{r,w}^h$ indicates warped random hair mask in I_{HA} as described in Section 3.3. The results show that all of our design decisions lead to an improvement of the FID score, successfully addressing both pose differences and occlusion inpainting.

Qualitative evaluations. Furthermore, we conduct a qualitative evaluation on the effect of SIM estimator. As mentioned in Section 3.3, SIM estimator guides our model to effectively inpaint occlusions by reflecting distinctive features of each region. According to Fig. 4, outputs of HairFIT with SIM estimator present more realistic inpainting quality compared to HairFIT without SIM estimator. In particular, while the output without SIM on the first row of Fig. 4 has blue artifacts on its clothes, the output with SIM does not. Moreover, unlike the output without SIM on the second row, the output with SIM has a clean forehead without brown artifacts.

5 Conclusion

We propose a two-stage pose-invariant hairstyle transfer model, HairFIT, which successfully transfers a target hairstyle to a source image when the source and the target have a different pose. In our model, a flow-based hair alignment network first aligns the target hairstyle with the source leveraging optical flow estimation. Then, a hair synthesis module generates output via an ALIAS generator with the help of a hair-agnostic image and a SIM estimator. Our SIM estimator guides the generator to inpaint occlusions in the source image which contain multiple semantic regions. The quantitative and qualitative results demonstrate the superiority of HairFIT over the existing methods.

Acknowledgement

This work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No. 2019-0-00075, Artificial Intelligence Graduate School Program (KAIST) and No. 2020-0-00368, A Neural-Symbolic Model for Knowledge Acquisition and Inference Techniques), the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIT) (No. NRF-2019R1A2C4070420), and Kakao Enterprise.

References

- [1] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *Proc. the International Conference on Learning Representations (ICLR)*, 2019.
- [2] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *Proc. of the IEEE international conference on computer vision (ICCV)*, 2017.
- [3] Seunghwan Choi, Sunghyun Park, Minsoo Lee, and Jaegul Choo. Viton-hd: High-resolution virtual try-on via misalignment-aware normalization. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2021.
- [4] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. v.d. Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *Proc. of the IEEE international conference on computer vision (ICCV)*, 2015.
- [5] Hang Gao, Huazhe Xu, Qi-Zhi Cai, Ruth Wang, Fisher Yu, and Trevor Darrell. Disentangling propagation and generation for video prediction. In *Proc. of the IEEE international conference on computer vision (ICCV)*, 2019.
- [6] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016.
- [7] Ke Gong, Yiming Gao, Xiaodan Liang, Xiaohui Shen, Meng Wang, and Liang Lin. Graphonomy: Universal human parsing via graph transfer learning. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2019.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [9] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

- [10] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Liteflownet: A lightweight convolutional neural network for optical flow estimation. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2018.
- [11] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2017.
- [12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. 2017.
- [13] Youngjoo Jo and Jongyoul Park. Sc-fegan: Face editing generative adversarial network with user’s sketch and color. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2019.
- [14] Taewoo Kim, Chaeyeon Chung, Sunghyun Park, Gyojung Gu, Keonmin Nam, Wonzo Choe, Jaesung Lee, and Jaegul Choo. K-hairstyle: A large-scale korean hairstyle dataset for virtual hair editing and hairstyle classification. In *Proc. of the IEEE International Conference on Image Processing (ICIP)*, 2021.
- [15] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2020.
- [16] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P Xing. Dual motion gan for future-flow embedded video prediction. In *Proc. of the IEEE international conference on computer vision (ICCV)*, 2017.
- [17] Ziwei Liu, Raymond A. Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *Proc. of the IEEE international conference on computer vision (ICCV)*, 2017.
- [18] A. Nagrani, J. S. Chung, and A. Zisserman. Voxceleb: a large-scale speaker identification dataset. In *INTERSPEECH*, 2017.
- [19] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *Proc. the International Conference on Learning Representations (ICLR)*, 2017.
- [20] Sunghyun Park, Kangyeol Kim, Junsoo Lee, Jaegul Choo, Joonseok Lee, Sookyoung Kim, and Yoonjae Choi. Vid-ode: Continuous-time video generation with neural ordinary differential equation. In *Proc. the AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [21] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2019.
- [22] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2019.

- [23] Tiziano Portenier, Qiyang Hu, Attila Szabó, Siavash Arjomand Bigdeli, Paolo Favaro, and Matthias Zwicker. Faceshop: Deep sketch-based face image editing. *ACM Transactions on Graphics (TOG)*, 37(4), 2018.
- [24] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *Proc. the International Conference on Learning Representations (ICLR)*, 2016.
- [25] Rohit Saha, Brendan Duke, Florian Shkurti, Graham Taylor, and Parham Aarabi. Loho: Latent optimization of hairstyles via orthogonalization. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2021.
- [26] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [27] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2019.
- [28] Zhentao Tan, Menglei Chai, Dongdong Chen, Jing Liao, Qi Chu, Lu Yuan, Sergey Tulyakov, and Nenghai Yu. Michigan: Multi-input-conditioned hair image generation for portrait editing. *ACM Transactions on Graphics (TOG)*, 39(4):1–13, 2020.
- [29] 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 *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2018.
- [30] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2021.
- [31] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [32] Shuai Yang, Zhangyang Wang, Jiaying Liu, and Zongming Guo. Deep plastic surgery: Robust and controllable image editing with human-drawn sketches. In *Proc. of the European Conference on Computer Vision (ECCV)*, 2020.
- [33] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proc. of the European Conference on Computer Vision (ECCV)*, 2018.
- [34] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *Proc. of the IEEE international conference on computer vision (ICCV)*, 2019.
- [35] Tao Yu, Zongyu Guo, Xin Jin, Shilin Wu, Zhibo Chen, Weiping Li, Zhizheng Zhang, and Sen Liu. Region normalization for image inpainting. In *Proc. the AAAI Conference on Artificial Intelligence (AAAI)*, 2020.

- [36] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 41(8):1947–1962, 2018.
- [37] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *Proc. the International Conference on Machine Learning (ICML)*, 2019.
- [38] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2018.