

What to Preserve and What to Transfer: Faithful, Identity-Preserving Diffusion-based Hairstyle Transfer

Chaeyeon Chung, Sunghyun Park, Jeongho Kim and Jaegul Choo

KAIST AI, South Korea

cy_chung@kaist.ac.kr, psh01087@gmail.com, {rlawjdghek, jchoo}@kaist.ac.kr

Abstract

Hairstyle transfer is a challenging task in the image editing field that modifies the hairstyle of a given face image while preserving its other appearance and background features. The existing hairstyle transfer approaches heavily rely on StyleGAN, which is pre-trained on cropped and aligned face images. Hence, they struggle to generalize under challenging conditions such as extreme variations of head poses or focal lengths. To address this issue, we propose a one-stage **hairstyle transfer diffusion** model, HairFusion, that applies to real-world scenarios. Specifically, we carefully design a hair-agnostic representation as the input of the model, where the original hair information is thoroughly eliminated. Next, we introduce a hair align cross-attention (Align-CA) to accurately align the reference hairstyle with the face image while considering the difference in their head poses. To enhance the preservation of the face image’s original features, we leverage adaptive hair blending during the inference, where the output’s hair regions are estimated by the cross-attention map in Align-CA and blended with non-hair areas of the face image. Our experimental results show that our method achieves state-of-the-art performance compared to the existing methods in preserving the integrity of both the transferred hairstyle and the surrounding features.

Introduction

With the recent advancement of generative models (Rombach et al. 2022; Saharia et al. 2022), image editing technologies (Zhang, Rao, and Agrawala 2023; Yang et al. 2023a; Hertz et al. 2022) have shown impressive results. The task of hairstyle transfer is one of the most challenging image editing tasks, focusing on modifying an input face image’s hairstyle to the reference hairstyle. This technology can offer users a preview of how different hairstyles would look on them, enhancing customer satisfaction. The key challenge of hairstyle transfer is to transfer the reference hairstyle to the face while preserving other features (*e.g.*, identity, clothing, and background) in the face image.

Recent hairstyle transfer (Kim et al. 2022; Khwanmuang et al. 2023; Wei et al. 2023; Nikolaev et al. 2024) allows users to manipulate hairstyles using StyleGAN2 (Karras et al. 2020), which can generate high-fidelity face images. However, such GAN-based approaches still suffer from several challenges. First, since most of them rely on StyleGAN2

pre-trained on FFHQ dataset (Karras, Laine, and Aila 2019), they often fail to generalize to face images with various head poses or focal lengths that lie outside the latent space of StyleGAN2. This hinders its application to real-world scenarios (Yang et al. 2022, 2023b) (see Fig. 5). Moreover, previous latent optimization methods have limitations in preserving fine-grained details of reference hairstyles (*e.g.*, curl), especially under extreme pose variations.

On the other hand, recent diffusion-based image editing demonstrates superior performance in manipulating human images compared to GANs. Thanks to the powerful generative prior of text-to-image diffusion models, there has been success in image editing based on human poses (Hu 2024; Kim et al. 2024b), depth (Zhang, Rao, and Agrawala 2023), reference images (Yang et al. 2023a), and clothing (Kim et al. 2024a; Choi et al. 2024). However, it is still under-explored to leverage diffusion models for hairstyle transfer.

To address the above-mentioned challenges, we propose a novel **hairstyle transfer diffusion** model called *HairFusion*, which generates a high-fidelity image with the reference hairstyle while faithfully preserving the surrounding features such as head shape, clothing, and backgrounds. We design a one-stage diffusion hairstyle transfer model building upon a pre-trained diffusion model (Rombach et al. 2022) to enhance generalizability on various face images. Inspired by recent literature on virtual try-on (Choi et al. 2021; Lee et al. 2022; Kim et al. 2024a), we posit that hairstyle transfer can be conceptualized as a process of filling in the hair region of a face image, conditioned on the reference hairstyle (*i.e.*, exemplar-based image inpainting). Adapting the diffusion-based image inpainting for the hairstyle transfer, we introduce a hair-agnostic representation and a hair align cross-attention (Align-CA). We first obtain a hair-agnostic representation by thoroughly eliminating the original hair information in the face image. Next, the Align-CA aligns the reference hairstyle with the hair region of the face image by learning the correspondence between them. Dense pose representations (Guler, Neverova, and Kokkinos 2018) are provided to the Align-CA to encourage the model to indicate the relative pose difference between reference hair and face.

Nevertheless, when editing the hairstyle via diffusion-based inpainting, challenges remain in effectively preserving essential original features such as identity, clothing, and background of the source face image. To solve this issue,



Figure 1: Generated results by HairFusion using in-the-wild images. Given a pair of a face image and a reference hairstyle image, our method can generate high-fidelity images. These results show our model’s generalizability to diverse face images.

we introduce adaptive hair blending. We exploit a cross-attention map to identify hair regions in the output and seamlessly blend them with non-hair areas of the face image during the inference. Our extensive experiments demonstrate that our method significantly outperforms the existing approaches in terms of synthesized image quality both qualitatively and quantitatively. Also, our method achieves superior performance than the existing diffusion models for exemplar-based image inpainting, even when applied to diverse, real-world images, as presented in Fig. 1. We summarize our contributions as follows:

- We present HairFusion, the first one-stage diffusion-based hairstyle transfer framework, capable of generalizing across arbitrary face images in real-world scenarios.
- We propose a hair align cross-attention module (Align-CA) that aligns the reference hairstyle with the face image by accounting for differences in head poses.
- Our novel adaptive hair blending effectively preserves the original features in non-hair regions by blending the estimated hair regions of the generated image with the non-hair regions of the face image during inference.

Related Work

Diffusion-based Image Editing. Text-to-image diffusion models (Saharia et al. 2022; Rombach et al. 2022) have demonstrated remarkable success in generating high-fidelity images based on textual descriptions. Several approaches have been proposed for image manipulation under diverse conditions, incorporating conditions such as pose (Hu 2024; Kim et al. 2024b), reference images (Yang et al. 2023a; Chen et al. 2024), and multiple conditions (Zhang, Rao, and Agrawala 2023; Kim et al. 2023) like depth maps and edge maps. Furthermore, several studies have extended to image inpainting based on the reference images, including objects (Yang et al. 2023a; Chen et al. 2024) or clothing (Kim

et al. 2024a). However, hairstyles, compared to other objects like clothing, involve more diverse inpainting regions due to the high variability in hair width and length across different styles. This study seeks to broaden the applicability of diffusion models by using them to transfer the hairstyle from the reference image to a source face image.

Hairstyle Transfer. The existing hairstyle transfer approaches exploit GANs to apply the reference hairstyle to the face image. MichiGAN (Tan et al. 2020) employed conditional generators incorporating hair attributes such as hair shape and appearance. LOHO (Saha et al. 2021) and Barber-shop (Zhu et al. 2021) introduced latent optimization techniques based on pre-trained StyleGAN (Karras et al. 2020) to preserve the overall structure of the face image while reflecting the reference hairstyle. To broaden the scope of hairstyle transfer applications, HairFIT (Chung et al. 2021), StyleYourHair (Kim et al. 2022), StyleGANSalon (Khwanmuang et al. 2023), and HairFastGAN (Nikolaev et al. 2024) have developed the pose-invariant hairstyle transfer, accommodating significant differences in head pose between the face and reference images. Moreover, HairCLIP (Wei et al. 2022) and HairCLIPv2 (Wei et al. 2023) interactively edit hairstyles based on multiple conditions, such as text descriptions, reference images, masks, and sketches. On the other hand, HairNeFR (Chang, Kim, and Kim 2023) leverages neural rendering to geometrically align target hair in the volumetric space. However, due to the capability of pre-trained StyleGAN, most previous methods still have limitations in generalizing to the face images with various head poses or focal length and preserving fine-grained details of reference hairstyles (*i.e.*, texture and curl). In this paper, we introduce a novel diffusion-based hairstyle transfer model performing on multi-view images, enhancing the applicability of hairstyle transfer in real-world scenarios.

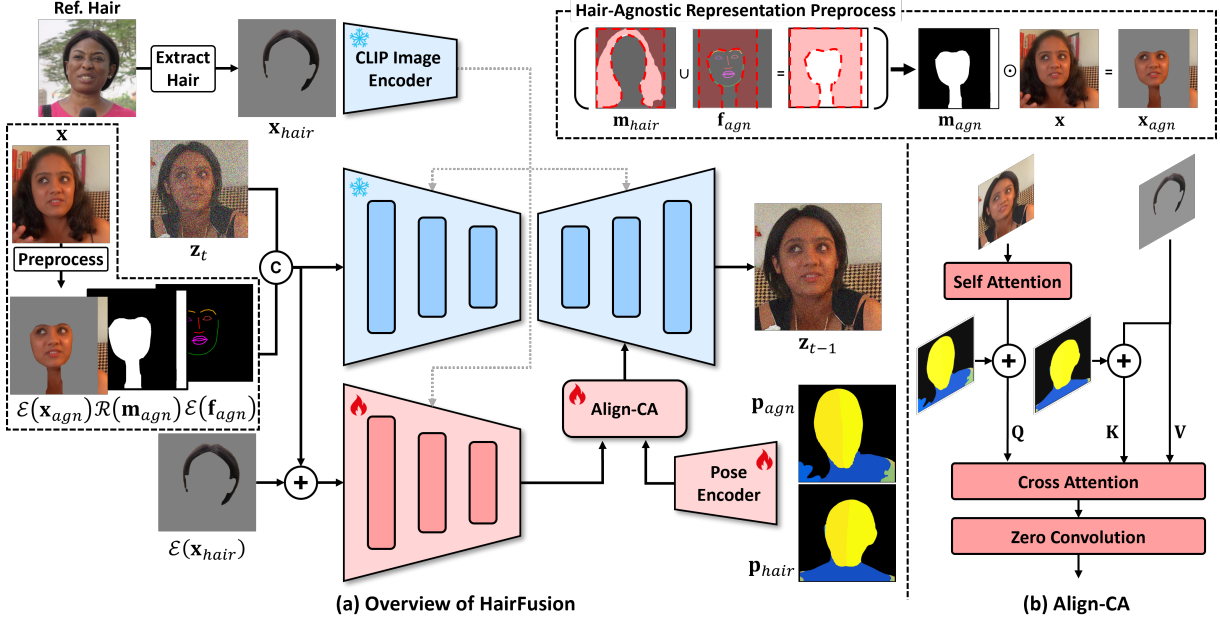


Figure 2: Overall pipeline of HairFusion. (a) HairFusion consists of a pre-trained U-Net, a hair encoder, a hair align cross-attention (Align-CA), and a pose encoder. We first preprocess hair-agnostic image \mathbf{x}_{agn} using a hair mask \mathbf{m}_{hair} and a face outline image \mathbf{f}_{agn} . Then, we provide \mathbf{x}_{agn} , a hair-agnostic mask \mathbf{m}_{agn} , a hair image \mathbf{x}_{hair} and dense pose images \mathbf{p}_{agn} , \mathbf{p}_{hair} as inputs to the model. (b) To inject \mathbf{x}_{hair} into \mathbf{x}_{agn} , we leverage the Align-CA, which aligns the hair features with the face features via cross-attention. Here, the pose features are added to the query (Q) and the key (K) as additional guidance.

Method

Preliminary

Large-scale diffusion probabilistic models have shown promising performance in image generation. Diffusion models (Ho, Jain, and Abbeel 2020) learn to generate images from a target data distribution by progressively denoising a normally distributed variable. Stable Diffusion (SD) (Rombach et al. 2022) is a latent diffusion model (LDM) that performs the denoising process in the latent space of an auto-encoder. To be specific, a pre-trained encoder (\mathcal{E}) first transforms an input image \mathbf{x} to latent feature $\mathbf{z}_0 = \mathcal{E}(\mathbf{x})$. Then, a forward diffusion process is performed with a pre-defined variance schedule β_t , following denoising diffusion probabilistic models (Ho, Jain, and Abbeel 2020):

$$q(\mathbf{z}_t|\mathbf{z}_0) = \mathcal{N}(\mathbf{z}_t; \sqrt{\bar{\alpha}_t}\mathbf{z}_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (1)$$

where $t \in \{1, \dots, T\}$ and T denotes the number of steps in the forward diffusion process. α_t is defined to be $1 - \beta_t$ and $\bar{\alpha}_t$ to be $\prod_{s=1}^t \alpha_s$. SD is trained with the following loss:

$$\mathcal{L}_{LDM} = \mathbb{E}_{\mathcal{E}(\mathbf{x}), \mathbf{y}, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(\mathbf{z}_t, t, \tau_\theta(\mathbf{y}))\|_2^2], \quad (2)$$

where $\epsilon_\theta(\cdot)$ is a denoising U-Net (Ronneberger, Fischer, and Brox 2015) and $\tau_\theta(\cdot)$ is a CLIP (Radford et al. 2021) text encoder that embeds the text prompt \mathbf{y} .

Overview

We address hairstyle transfer from the perspective of exemplar-based image inpainting. HairFusion aims to newly generate the masked hair region of a source face image

$\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$ according to the reference hairstyle $\mathbf{x}_{hair} \in \mathbb{R}^{3 \times H \times W}$. As presented in Fig. 2(a), we first preprocess a hair-agnostic image $\mathbf{x}_{agn} \in \mathbb{R}^{3 \times H \times W}$ that faithfully removes the original hair and potential hair regions for various reference hairstyles while preserving the regions that need to be reconstructed. Then, HairFusion takes \mathbf{x}_{agn} , \mathbf{x}_{hair} , and face outline image $\mathbf{f}_{agn} \in \mathbb{R}^{3 \times H \times W}$ as the input during denoising. Following Paint-by-Example (Yang et al. 2023a), we inject the class token of \mathbf{x}_{hair} extracted from a pre-trained CLIP image encoder into the denoising U-Net as an exemplar condition. Additionally, HairFusion leverages a hair encoder to extract detailed spatial features of \mathbf{x}_{hair} . We align the features of \mathbf{x}_{hair} with \mathbf{x}_{agn} via a hair align cross-attention (Align-CA). The Align-CA is designed to consider the relative difference of head poses, focal lengths, and face shape between \mathbf{x}_{agn} and \mathbf{x}_{hair} based on their dense pose features. We obtain dense pose features by injecting \mathbf{p}_{agn} and \mathbf{p}_{hair} into a pose encoder \mathcal{E}_p . During the inference, our newly proposed adaptive hair blending enhances the preservation of the original features in the source face image, which are largely occluded by \mathbf{m}_{agn} .

Hair-Agnostic Representation

To apply various reference hairstyles, all the potential hair regions in the source face \mathbf{x} need to be masked as an inpainting region. Also, the remaining hairstyle information in \mathbf{x} can harm the model’s generalization ability at inference. To address this issue, we design a hair-agnostic representation \mathbf{x}_{agn} that eliminates potential hair regions for \mathbf{x}_{hair} and the original hair of \mathbf{x} while preserving the regions that need

to be maintained. As illustrated in the top-right of Fig. 2, we first zero out the hair region of \mathbf{x} according to its hair segmentation mask \mathbf{m}_{hair} . Next, we remove the potential hair region based on the face outline image \mathbf{f}_{agn} . To be specific, we remove the area above the eyebrows and the region extending from the leftmost to the rightmost coordinate along the jawline with additional margins. We preserve the face region that needs to be accurately reconstructed. Note that we also preserve the neck and body beneath the chin, where hair typically does not exist. More details are illustrated in the supplementary material. Our hair-agnostic representation enables the model to process a wide variety of reference hairstyles while maintaining the core identity features.

Hair Align Cross-Attention

We leverage the cross-attention (CA) to align the reference hairstyle \mathbf{x}_{hair} with \mathbf{x}_{agn} . As shown in Fig. 2(b), our Align-CA exploits reference hairstyle features as the key (\mathbf{K}) and value (\mathbf{V}), where the output of the followed self-attention layer is given as the query (\mathbf{Q}). We extract intermediate features of the reference hairstyle using a trainable hair encoder whose initial weights are copied from the denoising U-Net.

Since \mathbf{x}_{hair} can vary in face shapes and head poses, it is challenging to accurately determine the hair shape or length based solely on \mathbf{x}_{hair} . For example, even if the hair in \mathbf{x}_{hair} appears to be long, it would be short hair if the reference’s face shape is elongated. Hence, we add the features of dense pose images, \mathbf{p}_{agn} and \mathbf{p}_{hair} , to \mathbf{Q} and \mathbf{K} in Align-CA:

$$\text{Align-CA} = \text{softmax}\left(\frac{(\mathbf{Q} + \mathcal{E}_p(\mathbf{p}_{agn})) \cdot (\mathbf{K} + \mathcal{E}_p(\mathbf{p}_{hair}))^T}{\sqrt{d}}\right) \cdot \mathbf{V}, \quad (3)$$

where \mathcal{E}_p is a pose encoder consisting of a stack of convolutional layers and d indicates the dimension of the feature. In this way, we can assist our model to accurately approximate the hair shape and length in the generated image in relation to the face based on the correspondence between \mathbf{x}_{hair} and \mathbf{x}_{agn} . We replace a few CA layers in the denoising U-Net decoder with our Align-CA.

Training Strategy

For training, we use a multi-view dataset to encourage the model to learn hair alignment across various head poses. Ideally, training samples consist of pairs with the same identity but different hairstyles and head poses. Since no existing datasets include such pairs, we utilize a multi-view dataset as an alternative to obtain pairs with the same identity and hairstyle but different head poses. During the training, one pair is given as \mathbf{x} , and the other’s hairstyle is given as \mathbf{x}_{hair} .

We train the hair encoder, Align-CA, and pose encoder (\mathcal{E}_p) of HairFusion with the following objective:

$$\mathcal{L}_{LDM} = \mathbb{E}_{\zeta, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(\zeta, t, \tau_\phi(\mathbf{x}_{hair}), \mathcal{E}(\mathbf{x}_{hair}))\|_2^2], \quad (4)$$

where

$$\zeta = ([\mathbf{z}_t; \mathcal{E}(\mathbf{x}_{agn}); \mathcal{R}(\mathbf{m}_{agn}); \mathcal{E}(\mathbf{f}_{agn})], \mathcal{E}_p(\mathbf{p}_{agn}), \mathcal{E}_p(\mathbf{p}_{hair})), \quad (5)$$

τ_ϕ is a CLIP image encoder, and \mathcal{R} indicates resize function.

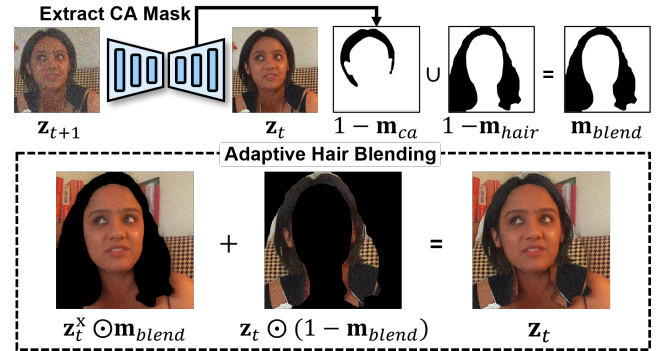


Figure 3: Overview of adaptive hair blending. We obtain \mathbf{m}_{blend} using \mathbf{m}_{ca} extracted from CA maps in Align-CA and the source hair mask \mathbf{m}_{hair} . \mathbf{m}_{blend} blends the generated hair features with the other features in the source.

Adaptive Hair Blending

One major challenge in hairstyle transfer is to preserve the original features in the source face image except for hair. However, the occluded region in \mathbf{x}_{agn} inevitably changes due to the autoencoder’s reconstruction error, degrading performance. As a naive solution, one can apply latent blending as in Blended Diffusion (Avrahami, Lischinski, and Fried 2022) using a hair-agnostic mask \mathbf{m}_{agn} as a blending mask. However, since \mathbf{m}_{agn} is designed to eliminate all the potential hair regions of the generated image, using \mathbf{m}_{agn} as a blending mask unnecessarily loses useful information of \mathbf{x} .

To address this issue, we propose an adaptive hair blending that maximally preserves the original features in \mathbf{x} . Since our Align-CA learns to align \mathbf{x}_{hair} with \mathbf{x}_{agn} , the aligned hair regions of the generated image are expected to be activated in the CA maps of Align-CA. With this motivation, we derive an adaptive blending mask (\mathbf{m}_{blend}) by leveraging the CA mask (\mathbf{m}_{ca}) which indicates the hair region of the output and the source hair mask (\mathbf{m}_{hair}) as illustrated in Fig. 3. Specifically, \mathbf{m}_{ca} is computed by normalizing and thresholding the CA maps extracted from the Align-CA layers. We obtain \mathbf{m}_{blend} with the union of $1 - \mathbf{m}_{ca}$ and $1 - \mathbf{m}_{hair}$, indicating the regions to preserve in \mathbf{x} as 1. Finally, we perform the adaptive hair blending with \mathbf{m}_{blend} , and update \mathbf{z}_t at timestep t with the blended feature before proceeding to the next timestep ($t - 1$) as follows:

$$\mathbf{z}_t \leftarrow \mathbf{z}_t^x \odot \mathbf{m}_{blend} + \mathbf{z}_t \odot (1 - \mathbf{m}_{blend}), \quad (6)$$

where \mathbf{z}_t^x indicates the noisy latent of \mathbf{x} at timestep t , and \mathbf{m}_{blend} is a binary mask, where the regions to be preserved are 1. We apply adaptive hair blending during the final N timesteps of the denoising process, and N is set as 10 in the entire experiment. In this way, we can faithfully preserve the features of the source face image \mathbf{x} , minimizing the reconstruction error of the autoencoder.

Experiment

Experimental Setup

Dataset. We utilize two multi-view datasets, the K-hairstyle (Kim et al. 2021) and the CelebV-Text (Yu et al.




Figure 4: Qualitative comparison with the diffusion-based baselines.



Figure 5: Qualitative comparison with baselines using web-crawled images.

2023) datasets for the experiments. The K-hairstyle dataset consists of 500,000 high-resolution face images with various focal lengths and head poses with more than 6,400 identities. Following the Style Your Hair (Kim et al. 2022), we excluded the images whose hairstyle or face is significantly occluded. Due to the privacy issue, we blur the face of the images from the K-hairstyle dataset. Also, the CelebV-Text dataset is a large-scale facial text-video dataset that contains 70,000 in-the-wild face video clips with text pairs. For the experiments, we randomly sampled 6,000 video clips with different identities and sampled 20 frames from each video. We remove images whose facial landmarks or hair masks are not detected. All the images are resized to 512×512 in the experiments. More details for dataset preprocessing are described in the supplementary material.

Baselines. We first compare our HairFusion with state-of-the-art diffusion models for exemplar-based image inpainting: Paint-by-Example (PbE) (Yang et al. 2023a) and Anydoor (Chen et al. 2024). For a fair comparison, we replace their input with our newly designed hair-agnostic representation during the training and inference. Since Anydoor uses various multi-view datasets to train the model, we fine-tune the Anydoor with our dataset instead of training the model from scratch. Also, we compare our method with StyleGAN-based hairstyle transfer approaches, Style Your Hair (SYH) (Kim et al. 2022) and HairCLIPv2 (Wei et al. 2023). We apply StyleGAN (Karras et al. 2020) trained on our dataset to StyleGAN-based approaches. Since the methods based on StyleGAN pre-trained with the FFHQ dataset (Karras, Laine, and Aila 2019) mainly tackle

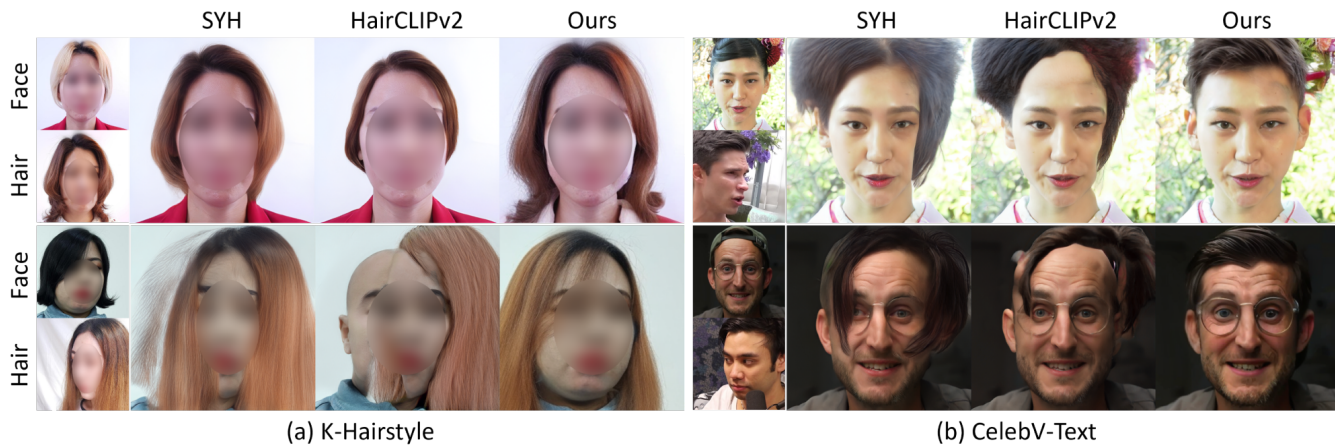


Figure 6: Qualitative comparison with StyleGAN-based methods using cropped and aligned images.

Dataset	Method	FID↓	SSIM↑	PSNR↑	LPIPS↓
K-Hairstyle	SYH	22.84	0.62	18.41	0.32
	PbE	13.26	0.59	19.26	0.29
	Anydoor	19.00	0.63	18.78	0.27
	Ours	10.82	0.70	21.15	0.18
CelebV-Text	SYH	35.77	0.62	21.14	0.30
	PbE	32.72	0.41	14.55	0.42
	Anydoor	35.70	0.50	15.45	0.37
	Ours	23.50	0.74	22.55	0.18

Table 1: Quantitative comparison to baselines, including SYH, PbE, and Anydoor.

cropped aligned faces, we also compare our method with StyleGAN-based methods using cropped aligned images. Specifically, we crop and align the test images in the way the FFHQ dataset is preprocessed for StyleGAN-based approaches. Also, we crop and align the results of HairFusion in the same way for comparison.

Evaluation Metric. We evaluate our method on two tasks, hairstyle transfer and reconstruction. Hairstyle transfer modifies hairstyle between two images with different identities and hairstyles. We measure the fr chet inception distance (FID) score (Heusel et al. 2017) to evaluate how similar the distributions of the synthesized and the real images are. In the reconstruction task, we utilize two images with the same identity and hairstyle but different head poses. One pair is considered the source face (*i.e.*, the ground truth image for the model to reconstruct) while the other serves as the reference hairstyle image. We measure SSIM, PSNR, and LPIPS for the reconstruction task to evaluate if the model accurately reflects the detailed features (*e.g.*, shape, length, color, etc.) of the reference hair in the result. For the evaluation, we randomly sample 2,000 pairs from the test sets.

Comparison to Baselines

Qualitative Comparison. Fig. 4 presents a comparison between our model and diffusion-based image inpainting models, PbE (Yang et al. 2023a) and Anydoor (Chen et al. 2024), using the K-Hairstyle and CelebV-Text datasets. PbE struggles to reflect the shape or color of the reference hair in the output. This is because PbE injects the reference hair fea-

Datset	Method	FID↓	SSIM↑	PSNR↑	LPIPS↓
K-Hairstyle	SYH	25.79	0.50	15.56	0.36
	HairCLIPv2	24.36	0.48	14.45	0.38
	Ours	15.41	0.55	19.26	0.30
CelebV-Text	SYH	35.69	0.65	22.09	0.28
	HairCLIPv2	35.40	0.66	21.21	0.29
	Ours	22.00	0.70	21.77	0.21

Table 2: Quantitative comparison to StyleGAN-based baselines using cropped and aligned dataset.

tures using only the class token extracted by the CLIP image encoder, which results in a significant loss of detailed hair shape and color information. While Anydoor preserves the features of the reference hair, it fails to maintain the head shape or clothing of the source face image.

Additionally, we conduct a comparison with web-crawled images in Fig. 5 using ours and the baselines trained on the K-Hairstyle dataset. Overall, SYH produces blurry outputs compared to the diffusion-based models. Although SYH employs StyleGAN trained on multi-view images, it fails to preserve the source face’s identity and loses the detail of the reference hair when applied to in-the-wild images containing diverse poses. Moreover, PbE and Anydoor struggle to maintain both the hairstyle features of the reference hair and the non-hair features in the face image.

Lastly, we compare our method to StyleGAN-based methods using cropped and aligned test sets in Fig. 6. The outputs generated by SYH and HairCLIPv2 have blurrier hair compared to ours. Also, SYH and HairCLIPv2 fail to align the reference hair to a face when they have a large pose difference, even when the inputs are cropped and aligned.

Unlike the existing methods, HairFusion generates realistic outputs not only in both the K-Hairstyle and CelebV-Text datasets but also in the web-crawled in-the-wild images while preserving both the fine details of the reference hair and the surrounding features in the source face image.

Quantitative Comparison. Table 1 shows that our model achieves superior performance compared to the existing methods, including SYH (Kim et al. 2022), PbE (Yang et al. 2023a), and Anydoor (Chen et al. 2024). For a fair comparison, we apply our newly designed hair-agnostic representa-




Figure 7: Visualization of estimated hair mask in Align-CA. The CA masks are obtained by normalizing and thresholding the CA maps in Align-CA. The CA masks successfully indicate the hair regions of the output. t indicates the timestep of the reverse denoising process, where the total timestep is 50.



Figure 8: Qualitative evaluation for the ablation study using K-Hairstyle dataset. The red box indicates the reference hair.

tion to the diffusion-based baselines. PbE achieves inferior performance in the reconstruction task (*i.e.*, SSIM, PSNR, and LPIPS), showing that it fails to reflect the detailed features of the reference hair. Although SYH and Anydoor achieve better reconstruction performance than PbE, they produce unrealistic outputs, achieving higher FID scores.

Also, we compare our method to the existing StyleGAN-based hairstyle transfer methods in Table 2 using cropped and aligned images. Our method achieves superior performance in both hairstyle transfer and the reconstruction task, except for CelebV-Text PSNR. Since most non-hair regions in face images are eliminated in cropped and aligned images, our adaptive hair blending may show only a slight improvement or comparable performance in the reconstruction task.

Ablation Study

We evaluate our method by gradually adding each component of HairFusion to the baseline. As in Fig. 8 and Table 3, we start from the baseline which only contains the denoising U-Net, CLIP image encoder, and the hair encoder. Then, we gradually add Align-CA and the adaptive hair blending. The results show that the baseline struggles to estimate the exact hair shape and length of the aligned reference hair, achieving lower reconstruction performance. Although adding Align-CA largely improves hair alignment performance, it still fails to preserve detailed features in the non-hair region of

	FID↓	SSIM↑	PSNR↑	LPIPS↓
Baseline	15.43	0.58	18.09	0.30
+ Align-CA	15.59	0.62	19.33	0.27
+ Adaptive Hair Blending	10.82	0.70	21.15	0.18

Table 3: Quantitative ablation study using K-Hairstyle.

the face image such as clothing. Our method with adaptive hair blending achieves the best performance in both hairstyle transfer and reconstruction. Fig. 8 shows that only ours successfully maintains the reference hair features as well as the clothing’s details of the face image.

Analysis of Adaptive Hair Blending

Adaptive hair blending leverages the CA mask \mathbf{m}_{ca} in the Align-CA to estimate hair regions of the output. Fig. 7 visualizes \mathbf{m}_{ca} extracted from Align-CA every 10 timesteps as the 50 timesteps of the reverse denoising process proceed. The first and the second row present a test pair from K-Hairstyle and CelebV-Text, respectively. We visualize \mathbf{m}_{ca} from the 7-th CA map in Align-CA while we use the 6-th and 7-th CA map in adaptive hair blending. The figure shows that \mathbf{m}_{ca} successfully indicates the hair regions of the generated image. HairFusion effectively reconstructs the original features of \mathbf{x} by blending the generated hair features and the non-hair region features in \mathbf{x} .

Conclusion

This paper proposes the first one-stage diffusion-based hairstyle transfer model, HairFusion, conceptualizing hairstyle transfer as an exemplar-based image inpainting. HairFusion introduces Align-CA that aligns the target hairstyle with a face image based on dense pose features, accounting for their pose difference. Our novel adaptive hair blending technique allows HairFusion to blend the transferred reference hair features with the source face’s other appearance and background features based on the hair region estimated by CA maps of Align-CA. HairFusion achieves state-of-the-art performance compared to the existing approaches, including StyleGAN-based methods and diffusion models for exemplar-based inpainting. We also demonstrate that HairFusion can generalize to in-the-wild samples with diverse head poses and focal lengths.

Acknowledgements

This work was supported by the Institute of Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIT) (No.RS-2019-II190075 Artificial Intelligence Graduate School Program(KAIST) and RS-2021-II212068, Artificial Intelligence Innovation Hub) and the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT)(No. 2022R1A5A7083908).

References

- Avrahami, O.; Lischinski, D.; and Fried, O. 2022. Blended Diffusion for Text-Driven Editing of Natural Images. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 18208–18218.
- Chang, S.; Kim, G.; and Kim, H. 2023. HairNeRF: Geometry-Aware Image Synthesis for Hairstyle Transfer. In *Proc. of the IEEE international conference on computer vision (ICCV)*, 2448–2458.
- Chen, X.; Huang, L.; Liu, Y.; Shen, Y.; Zhao, D.; and Zhao, H. 2024. Anydoor: Zero-shot object-level image customization. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 6593–6602.
- Choi, S.; Park, S.; Lee, M.; and Choo, J. 2021. VITON-HD: High-resolution virtual try-on via misalignment-aware normalization. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*.
- Choi, Y.; Kwak, S.; Lee, K.; Choi, H.; and Shin, J. 2024. Improving diffusion models for virtual try-on. In *Proc. of the European Conference on Computer Vision (ECCV)*.
- Chung, C.; Kim, T.; Nam, H.; Choi, S.; Gu, G.; Park, S.; and Choo, J. 2021. HairFIT: Pose-Invariant Hairstyle Transfer via Flow-based Hair Alignment and Semantic-Region-Aware inpainting. In *Proc. of the British Machine Vision Conference (BMVC)*. British Machine Vision Association.
- Guler, R. A.; Neverova, N.; and Kokkinos, I. 2018. DensePose: Dense Human Pose Estimation In The Wild. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*, 33: 6840–6851.
- Hu, L. 2024. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 8153–8163.
- Karras, T.; Laine, S.; and Aila, T. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*.
- Khwanmuang, S.; Phongthawee, P.; Sangkloy, P.; and Suwajanakorn, S. 2023. StyleGAN Salon: Multi-View Latent Optimization for Pose-Invariant Hairstyle Transfer. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 8609–8618.
- Kim, J.; Gu, G.; Park, M.; Park, S.; and Choo, J. 2024a. Stableviton: Learning semantic correspondence with latent diffusion model for virtual try-on. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 8176–8185.
- Kim, J.; Kim, M.-J.; Lee, J.; and Choo, J. 2024b. TCAN: Animating Human Images with Temporally Consistent Pose Guidance using Diffusion Models. *Proc. of the European Conference on Computer Vision (ECCV)*.
- Kim, K.; Park, S.; Lee, J.; and Choo, J. 2023. Reference-based image composition with sketch via structure-aware diffusion model. In *Proc. of the IEEE conference on computer vision and pattern recognition workshop (CVPRW)*.
- Kim, T.; Chung, C.; Kim, Y.; Park, S.; Kim, K.; and Choo, J. 2022. Style your hair: Latent optimization for pose-invariant hairstyle transfer via local-style-aware hair alignment. In *Proc. of the European Conference on Computer Vision (ECCV)*, 188–203. Springer.
- Kim, T.; Chung, C.; Park, S.; Gu, G.; Nam, K.; Choe, W.; Lee, J.; and Choo, J. 2021. 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)*, 1299–1303. IEEE.
- Lee, S.; Gu, G.; Park, S.; Choi, S.; and Choo, J. 2022. High-resolution virtual try-on with misalignment and occlusion-handled conditions. In *Proc. of the European Conference on Computer Vision (ECCV)*, 204–219. Springer.
- Nikolaev, M.; Kuznetsov, M.; Vetrov, D.; and Alanov, A. 2024. HairFastGAN: Realistic and Robust Hair Transfer with a Fast Encoder-Based Approach. *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *Proc. the International Conference on Machine Learning (ICML)*, 8748–8763. PMLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, 234–241. Springer.

Saha, R.; Duke, B.; Shkurti, F.; Taylor, G.; and Aarabi, P. 2021. LOHO: Latent optimization of hairstyles via orthogonalization. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*.

Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*, 35: 36479–36494.

Tan, Z.; Chai, M.; Chen, D.; Liao, J.; Chu, Q.; Yuan, L.; Tulyakov, S.; and Yu, N. 2020. MichiGAN: Multi-input-conditioned hair image generation for portrait editing. *Proc. the ACM Transactions on Graphics (ToG)*, 39(4): 1–13.

Wei, T.; Chen, D.; Zhou, W.; Liao, J.; Tan, Z.; Yuan, L.; Zhang, W.; and Yu, N. 2022. Hairclip: Design your hair by text and reference image. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 18072–18081.

Wei, T.; Chen, D.; Zhou, W.; Liao, J.; Zhang, W.; Hua, G.; and Yu, N. 2023. HairCLIPv2: Unifying Hair Editing via Proxy Feature Blending. In *Proc. of the IEEE international conference on computer vision (ICCV)*, 23589–23599.

Yang, B.; Gu, S.; Zhang, B.; Zhang, T.; Chen, X.; Sun, X.; Chen, D.; and Wen, F. 2023a. Paint by example: Exemplar-based image editing with diffusion models. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 18381–18391.

Yang, S.; Jiang, L.; Liu, Z.; ; and Loy, C. C. 2023b. StyleGANEX: StyleGAN-Based Manipulation Beyond Cropped Aligned Faces. In *Proc. of the IEEE international conference on computer vision (ICCV)*.

Yang, S.; Jiang, L.; Liu, Z.; and Loy, C. C. 2022. VToonify: Controllable High-Resolution Portrait Video Style Transfer. *Proc. of the ACM SIGGRAPH Asia*.

Yu, J.; Zhu, H.; Jiang, L.; Loy, C. C.; Cai, W.; and Wu, W. 2023. Celebv-text: A large-scale facial text-video dataset. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 14805–14814.

Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proc. of the IEEE international conference on computer vision (ICCV)*, 3836–3847.

Zhu, P.; Abdal, R.; Femiani, J.; and Wonka, P. 2021. Barber-shop: GAN-based Image Compositing using Segmentation Masks. *Proc. of the ACM SIGGRAPH Asia*, 40(6).