

Stable-Hair: Real-World Hair Transfer via Diffusion Model

Yuxuan Zhang¹, Qing Zhang³, Yiren Song⁴, Jichao Zhang⁵, Hao Tang⁶, Jiaming Liu^{2*}

¹Shanghai Jiao Tong University

²Tiamat AI

³Shenyang Institute of Automation Chinese Academy of Sciences

⁴National University of Singapore

⁵Ocean University of China

⁶Peking University

zyx153@sjtu.edu.cn, zhangqing221@mailsucas.ac.cn, yiren@nus.edu.sg,
zhang163220@gmail.com, haotang@pku.edu.cn, jmliu1217@gmail.com

Abstract

Current hair transfer methods struggle to handle diverse and intricate hairstyles, limiting their applicability in real-world scenarios. In this paper, we propose a novel diffusion-based hair transfer framework, named *Stable-Hair*, which robustly transfers a wide range of real-world hairstyles to user-provided faces for virtual hair try-on. To achieve this goal, our *Stable-Hair* framework is designed as a two-stage pipeline. In the first stage, we train a Bald Converter alongside stable diffusion to remove hair from the user-provided face images, resulting in bald images. In the second stage, we specifically designed a Hair Extractor and a Latent IdentityNet to transfer the target hairstyle with highly detailed and high-fidelity to the bald image. The Hair Extractor is trained to encode reference images with the desired hairstyles, while the Latent IdentityNet ensures consistency in identity and background. To minimize color deviations between source images and transfer results, we introduce a novel Latent ControlNet architecture, which functions as both the Bald Converter and Latent IdentityNet. After training on our curated triplet dataset, our method accurately transfers highly detailed and high-fidelity hairstyles to the source images. Extensive experiments demonstrate that our approach achieves state-of-the-art performance compared to existing hair transfer methods.

Introduction

Hair transfer is one of the most challenging tasks in the virtual try-on domain. The objective of this task is to transfer hair color, shape, and structure attributes from a reference image to a user-provided source image while preserving the identity and background of the source image. In recent years, advances in Generative Adversarial Networks (GANs) (Tan et al. 2020; Zhang and Zheng 2018; Guo et al. 2022; Zhu et al. 2022; Chang, Kim, and Kim 2023; Chung et al. 2022; Zhu et al. 2021; Saha et al. 2021; Wei et al. 2022, 2023; Nikolaev et al. 2024; Khwanmuang et al. 2023; Kim et al. 2022) have driven significant progress in this field. However, these GAN-based methods often struggle to handle the diverse and complex hairstyles encountered in real-world scenarios, severely limiting their effectiveness in practical applications. Recently, diffusion models have emerged as

SOTA methods in the field of image generation. These models not only enable more stable training but also demonstrate impressive results in terms of diversity and fidelity. Consequently, we are intrigued by the question: “*Can we leverage the powerful capabilities of diffusion models to achieve more stable and high-precision hair transfer?*”

Building on the aforementioned considerations, in this paper, we propose a novel hair transfer framework based on diffusion models, named *Stable-Hair*. Inspired by method (Wei et al. 2023) which converts different editing conditions (e.g., text, sketch) into different proxies in the StyleGAN $\mathcal{W}+$ space, we designed a two-stage paradigm to ensure the precision and naturalness of our hair transfer process. In the first stage, we train a Bald Converter to transform the user-provided source image into a bald proxy image. Furthermore, a key advancement in our methodology is the implementation of an automated data generation pipeline, which generates triplets for training. This pipeline utilizes the Large Language Model for generating text prompts, and the Stable Diffusion Inpainting model for creating reference images. This synthetic training data ensures the effective training of our framework, enabling it to robustly handle challenging hairstyles with remarkable fidelity. In the second stage, a Hair Extractor is specially trained to capture and encode hairstyles from reference images with unprecedented levels of detail and texture. A Latent IdentityNet is used to encode the source image, ensuring that the identity and background of the source image remain consistent throughout the transfer process. By integrating these two components with diffusion U-Net, we can effectively inject the features of the reference hair into the source image. This approach not only allows the hair to adapt accurately to the new environment of the source image but also ensures that it is styled appropriately, appearing natural and matching the original appearance of the subject.

To rigorously preserve the non-hair regions and maintain the facial identity of the subject, we have developed an innovative architecture named Latent ControlNet. This structure integrates our Bald Converter and Latent IdentityNet. By transforming the pixel-space transfer process to operate within the latent space of the diffusion model, Latent ControlNet ensures color consistency in the non-hair regions of the source image during the transfer process.

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Stable-Hair is the first diffusion-based method for hairstyle transfer, capable of handling an extensive range of real-world hairstyles with exceptional robustness. Unlike previous methods, which often struggle with complex or intricate styles, Stable-Hair achieves remarkably detailed and high-fidelity transfers while preserving the original identity content.

Through extensive experiments, Stable-Hair has demonstrated its superior performance, significantly surpassing existing state-of-the-art hair transfer methods in terms of fidelity and fine-grained detail. Our approach sets a new standard for hair transfer technology, promising to revolutionize the virtual hair try-on experience with its advanced capabilities and innovative design.

In summary, our contributions are:

- In this paper, we propose the first diffusion-based hair style transfer framework, named *Stable-Hair*. Experiments demonstrate that our method significantly surpasses existing SOTA hair transfer methods in terms of fidelity and fine-grained detail.
- We utilize a Hair Extractor to encode reference images and inject detailed hair features. To ensure better source content consistency during the transfer process, we introduce a novel Latent ControlNet architecture. This architecture is utilized as a Bald Converter and a Latent IdentityNet, which maps the hairstyle transfer process from the pixel space to the latent space.
- We utilized our designed pipeline to process both video and image data. This pipeline plays a crucial role in the effective training of our framework by generating a high-quality dataset.

Related Works

Hair Style Transfer

The development of GAN-based methods (Tan et al. 2020; Zhang and Zheng 2018; Guo et al. 2022; Zhu et al. 2022; Chang, Kim, and Kim 2023; Chung et al. 2022; Zhu et al. 2021; Saha et al. 2021; Wei et al. 2022, 2023; Nikolaev et al. 2024; Khwanmuang et al. 2023; Kim et al. 2022; Shu et al. 2022) has significantly advanced the field of hairstyle transfer, with most current methods relying on GAN-based approaches.

In hair transfer, MichiGAN (Tan et al. 2020) decomposes hair into four orthogonal attributes and designs the corresponding modules to represent, process, and convert user input, which is then integrated to realize an end-to-end network. Barbershop (Zhu et al. 2021) proposes a novel latent space for image blending that excels at preserving detail and encoding spatial information, extracting hairstyle information from multiple reference images. Furthermore, LOHO (Saha et al. 2021) employs an optimization-based approach using GAN inversion to infill missing hair structure details in the latent space during hairstyle transfer, introducing two-stage optimization and gradient orthogonalization to enable disentangled latent space optimization of hair attributes. Hairmapper (Wu, Yang, and Jin 2022) trains a network to direct hair removal in StyleGAN’s latent space. Although these early methods can transfer simple hairstyles, they do not address the challenge of pose-invariant hairstyle transfer. Therefore, SYH (Kim et al. 2022) proposes a pose-invariant model that handles significant pose differences while maintaining local textures. Hairclip (Wei et al. 2022) utilizes CLIP for unified hair editing, with HairCLIPv2 (Wei et al. 2023) converting hair editing tasks into transfer tasks with varied proxies. StyleGAN-Salon (Khwanmuang et al. 2023) employs multi-view optimization and guide images to enhance hairstyle transfer accuracy. HairFastGAN (Nikolaev et al. 2024) addresses the challenge of hairstyle transfers in different poses and proposes a new encoder capable of generating images promptly.

Despite these advancements, GAN-based models have limitations in addressing complex hairstyle transfers in real-world scenarios. To overcome these deficiencies, we propose the first diffusion-based method, which generates high-quality and robust images, thereby setting a new benchmark in the field.

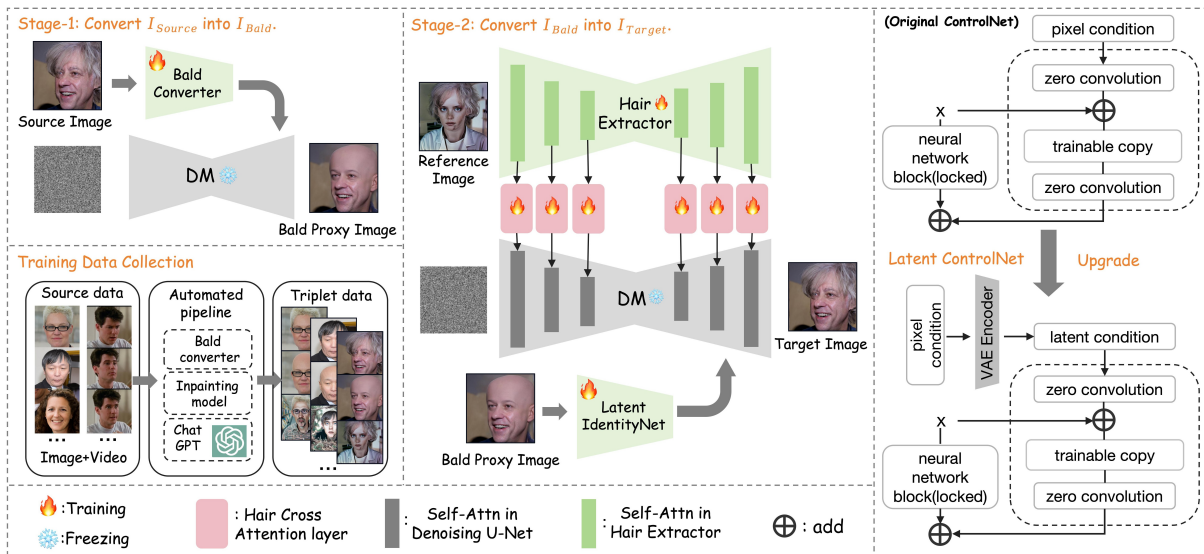


Figure 2: Overall schematics of our method. Our pipeline consists of two stages. First, the user’s input source image is transformed into a bald proxy image by utilizing a Bald Converter. In the second stage, we employ the pre-trained SD model along with a Hair Extractor to transfer the reference hair onto the bald proxy image. The Hair Extractor is responsible for capturing the intricate details and features of the reference hair. These features are then injected into the SD model through newly added hair cross-attention layers. After training on the triplet dataset constructed using our specially designed automated data pipeline, our method achieves highly detailed and high-fidelity hair transfers, resulting in natural and visually appealing outcomes.

Diffusion Models

Currently, diffusion models have attracted a great deal of attention and have seen significant advancements. As the most prominent generative models today, diffusion models have achieved state-of-the-art results across various image generation tasks, including text-to-image generation (Ramesh et al. 2022; Podell et al. 2023; Saharia et al. 2022a; Rombach et al. 2022; Saharia et al. 2022b), image editing (Chefer et al. 2023; Li et al. 2023; Epstein et al. 2023; Cao et al. 2023; Xie et al. 2023; Mou et al. 2023a; Bodur et al. 2023; Zhang et al. 2023; Tsaban and Passos 2023; Yang et al. 2024; Li et al. 2024), controllable generation (Zhang and Agrawala 2023; Mou et al. 2023b; Zhao et al. 2023; Ma et al. 2023), personalized image generation (Zhang et al. 2024b; Hu et al. 2022; Arar et al. 2023; Jia et al. 2023; Gal et al. 2022; Ruiz et al. 2023, 2022; Zhang et al. 2024c) and so on. In particular, the recent emergence of diffusion-based virtual try-on applications further demonstrates the formidable generative capabilities of diffusion models, enabling commercial grade virtual try-ons (Xu et al. 2024; Wang et al. 2024; Kim et al. 2023; Zeng et al. 2023) and virtual makeovers (Zhang et al. 2024d), which were previously unattainable with traditional GAN methods.

Diffusion models excel in generating high-quality images, especially when trained on extensive datasets. In this paper, We use diffusion models to achieve high-fidelity and robust hairstyle transfers, addressing the limitations of GAN-based methods and significantly enhancing the quality and realism of complex hairstyles.

Methodology

Overview

As shown in Fig. 2, our design divides the hairstyle transfer process into two stages. Firstly, the user’s input source image is transformed into a bald proxy image by using the Bald Converter. Secondly, our model is used to transfer the reference hair onto the bald proxy image. This two-stage approach ensures optimal stability in hairstyle transfer and maintains the source image content.

Specifically, given a reference image containing the desired hairstyle and a source image provided by users, Hair Extractor can stably and precisely encode a variety of reference hairstyles in real-world scenes and extract detailed hair features. These detailed hair features are fed to the diffusion U-Net structure. The Latent IdentityNet modules are used to process the bald proxy image that is converted from the source image by the Bald Converter.

Stage1: Standardizing to Bald State

Bald Converter. In our method, the source image is processed in two steps, the first step is to convert the source image into a bald proxy image by a *Bald Converter*. Specifically, we use the vanilla diffusion model in conjunction with our designed Latent ControlNet structure (introduced in the following section) to transfer the source image into the bald state.

After training, the Bald Converter achieves complete hair removal without requiring image alignment or cropping. This results in a cleaner, more consistent baseline for subsequent hairstyle transfer, leading to improved visual fidelity

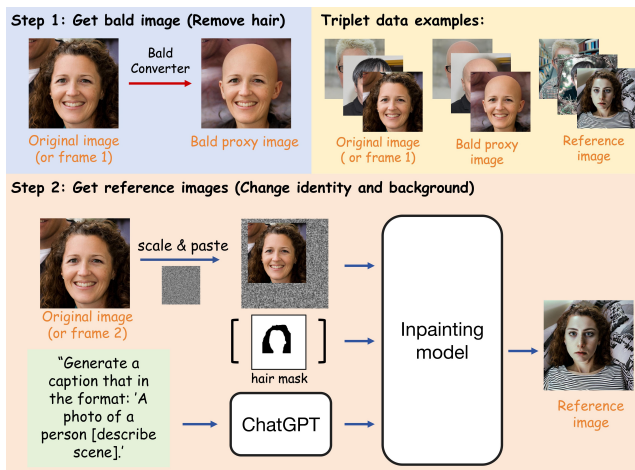


Figure 3: **Synthetic Training Data:** We propose an automated data generation pipeline to generate {Original image (or frame 1), Reference image, Bald proxy image} triplets for training. The pipeline uses ChatGPT to generate text prompts, the Stable Diffusion Inpainting model to generate reference images, and our pre-trained Bald converter to convert the original image or one of the frames sampled from videos into the bald proxy image.

and more realistic outcomes. Notably, standardizing source images to a bald state is crucial for enhancing model performance in hairstyle transfer. This process allows the model to concentrate more effectively on facial features and details by removing hair-related variability. Such standardization is instrumental in improving both training stability and transfer accuracy. The significance and impact of this approach are discussed in greater detail in the supplementary materials.

Training Data Collection

After training our Bald Converter, we can utilize the Bald Converter alongside the Stable Diffusion Inpainting model and ChatGPT to generate the triplet training data necessary to train the second-stage models.

Automatic Pipeline. As shown in Fig. 3, our pipeline to create the hairstyle pairing dataset involves two main steps. First, we use the bald converter to generate bald images, which serve as proxy images for bald heads during training. Second, based on the hairstyle masks from the original images, we use a Stable Diffusion inpainting model and ChatGPT to edit the non-hairstyle portions of the original dataset, altering identities and backgrounds to create the reference images for training. Ultimately, we obtain a triplet dataset consisting of the original images, the reference images, and the bald proxy images.

Video Strategy. Relying solely on image data, the reference and source images typically share the same head pose. This limitation prevents the model from effectively handling cases where the head poses in the reference and source images are misaligned. Video data, on the other hand, often includes frames of the same individual (with consistent

hairstyle) from different angles. By using video data to generate our dataset—selecting one frame as the source image and another with a different pose as the reference image—we can elegantly address this issue. This approach enhances the model’s robustness to pose variations.

In terms of image data, we used every single image as original images to generate reference images and bald proxy images, resulting in a total of 60,000 images. For video data, we sampled two frames from the same video. One frame was used to generate bald proxy images, and another frame was used to generate reference images. Resulting in a total of 90,000 images.

Stage2: Hair Transfer and Integration

Latent IdentityNet. In our method, the second step is to transfer the reference hair onto the bald proxy image. The maintenance of content consistency in the source image is crucial through our two steps. Any deviation from the intended goal, such as alterations in color or identity, would result in a content-inconsistent final image. Therefore, designing a maintenance module becomes a pivotal aspect of our hair transfer framework.

A baseline approach is using the ControlNet structure as a *Bald Converter* and *Latent IdentityNet* to ensure content consistency. However, our experimental findings indicate that while ControlNet effectively maintains the structural consistency of the source image, it struggles to preserve color consistency. As shown in Fig. 6, due to accumulated color deviations over these two steps, there are noticeable changes in the final colors. Why does the controlnet have a color difference? We assumed that the reason is that the pixel space in ControlNet and the latent space in U-Net represent image information in fundamentally different ways. The pixel space deals with the raw pixel values of an image, while the latent space involves a more abstract, high-dimensional representation created by the VAE encoder. For diffusion models, the misalignment between the characteristics and distribution of information in these two spaces can lead to difficulties in maintaining color consistency during the hair transfer process.

Therefore, as shown in Fig. 2 (right), we improve the ControlNet structure and propose a new variant named **Latent ControlNet**. Before the image is input into the ControlNet, the image is encoded into the latent space by the VAE encoder and then sent to the trainable copy of U-Net after a new trainable convolutional layer. Finally, we train our *Bald Converter* and *Latent IdentityNet* based on our proposed Latent ControlNet structure and get the best consistency practice.

Hair Extractor. Hair transfer needs to transfer the hair in the reference image in a detailed, complete, and accurate way, and our Hair Extractor is designed to achieve such transfer. Inspired by recent works (Wang et al. 2024; Zhang et al. 2024a) on reference image-guided generation, we utilize a trainable copy of the U-Net from pre-trained diffusion models as our Hair Extractor. Specifically, we encode our reference image through the Hair Extractor and collect the features of the self-attention layer in each transformer



Figure 4: Qualitative comparison of different methods. Compared to other approaches, our method achieves more refined and stable hairstyle transfer without the need for precise facial alignment or explicit masks for supervision.

block as detailed hair features. The features are then subsequently injected into the diffusion U-Net through newly added hair cross-attention layers. In each transformer block of the U-Net, we retain the original self-attention layers and add hair cross-attention layers. Detailed hair features are entered into the hair cross-attention layers and serve as the K(key) and V(value) features. Both attention layers in the U-Net share the Q feature. Finally, we simply add the output of hair cross-attention to the output of original self-attention.

Model Training

In the first stage, we train the bald converter using a straightforward approach similar to ControlNet on an existing dataset. This allows us to achieve a highly effective bald converter. In the second stage, we focus on training the Hair Extractor and the Latent IdentityNet. In these two training processes, we have used a variety of augmentations, which are crucial for adapting to real-world scenarios and achieving successful hair transfer. These augmentations include synchronized affine transformations applied to the source im-

age, the bald proxy image, and the target image.

The loss functions for both stages can be mathematically represented as follows:

$$L(\theta) := \mathbb{E}_{\mathbf{x}_0, t, \epsilon} \left[\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t, \mathbf{c}_s, \mathbf{c}_r)\|_2^2 \right], \quad (1)$$

where \mathbf{x}_t is a noisy image latent constructed by adding noise $\epsilon \in \mathcal{N}(\mathbf{0}, \mathbf{1})$ to the image latent \mathbf{x}_0 and the network $\epsilon_\theta(\cdot)$ is trained to predict the added noise, \mathbf{c}_s , \mathbf{c}_r represent the source condition input(source image or bald proxy image), and reference condition input, respectively. When training the bald converter, \mathbf{c}_r is None.

Experiments

Implementation Details

We employed Stable Diffusion V1-5 as the pre-trained diffusion model. The model training process utilized a two-stage approach. In the first stage, we trained the Bald Converter using the Non-Hair FFHQ dataset (Wu, Yang, and Jin 2022). This training was conducted on a single H800 GPU with a

Method	CLIP-I \uparrow	FID \downarrow	PSNR \uparrow	SSIM \uparrow	IDS \uparrow
Barbershop	0.431	46.178	30.397	0.629	0.760
HairFastGAN	0.426	<u>36.205</u>	30.383	<u>0.666</u>	0.762
SYH	0.429	41.543	30.489	0.645	0.712
Hairclip	0.391	44.359	28.855	0.615	0.697
Hairclip v2	0.419	37.456	31.552	0.642	<u>0.769</u>
Ours	0.434	35.128	<u>30.980</u>	0.680	0.779

Table 1: **Quantitative comparison** of different methods. Metrics that are bold and underlined represent methods that rank 1st and 2nd, respectively.

batch size of 16 and a learning rate of $5e-5$, over a total of 8,000 steps. In the second stage, we trained our Hair Extractor and Latent IdentityNet using the prepared triplet data. This stage was performed on 8 H800 GPUs with a batch size of 8 and a learning rate of $5e-5$, over a total of 100,000 steps. During inference, we followed the same two-stage approach as used in training. Both stages utilized the DDIM sampler with 30 sampling steps and the classifier-free guidance scale setting of 1.5.

Evaluation Metrics

Given a hairstyle reference image, the purpose of hair transfer is to apply the corresponding hairstyle and hair color attributes to the input image. We compare our method with current state-of-the-art methods: *Barbershop* (Zhu et al. 2021), *SYH* (Kim et al. 2022), *HairFastGAN* (Nikolaev et al. 2024), *Hairclip* (Wei et al. 2022), and *Hairclip v2* (Wei et al. 2023). All comparison algorithms use the default parameters from their official implementations.

To provide a thorough and objective assessment of the performance of each algorithm in different aspects of hairstyle transfer, we calculated the *FID* (Heusel et al. 2017) metrics for the source image and the generated target image. After hairstyle transfer, the identity and background information of the source image and the generated image should be consistent, so we use *SSIM* (Wang et al. 2004) and *PSNR* to evaluate the identity and background similarity between the source image and the generated target image. Notably, *PSNR* and *SSIM* are calculated at the intersected non-hair regions before and after editing. We also use *Insightface* (Deng et al. 2019) to evaluate *identity similarity (IDS)* between the original source image and the generated target image. To evaluate hairstyle transfer capabilities, we use *CLIP-I* (Radford et al. 2021) as a metric, which calculates the cosine similarity between the image embeddings of the reference image and the transferred hairstyle image.

Experiment Results

Qualitative Comparison. As illustrated in Fig. 4, we conducted qualitative comparison experiments on a variety of hairstyles. Overall, our method significantly outperforms other approaches in terms of the refinement and completeness of hairstyle transfer, while also maintaining the background and identity consistency of the source image to a



Figure 5: Visual comparison of hair removal using HairMapper. Our bald converter demonstrates robust performance, effectively converting source images across a wide range of poses, half-body shots, and even animated characters. In contrast, HairMapper struggles with these diverse scenarios, failing to maintain identity and background consistency.

great extent. Among the methods we compared, *Barbershop* tends to produce chaotic results for complex hairstyle transfers, as observed in the second and fifth rows. *Style Your Hair* and *HairFastGAN* perform relatively coarsely in hairstyle transfer, often neglecting the fine details of hair texture and color. *Hairclip* and *Hairclip v2* demonstrate the weakest ability in hairstyle transfer, struggling to accurately transfer the reference hairstyle. In contrast, our method consistently exhibits robust and stable transfer capabilities across different hairstyle styles and colors. Furthermore, We also compared the hair removal capabilities of our Bald Converter with *HairMapper* in Fig. 5. Our Bald Converter demonstrates robust and reliable performance, effectively transforming source images across various challenging scenarios. In contrast, *HairMapper* frequently struggles in these contexts, leading to inconsistencies in both identity and background.

Quantitative Comparison. The experiment uses the CelebA-HQ dataset (Karras et al. 2017) as experimental data, 2500 face images are randomly selected as input from the CelebA-HQ dataset with an equal number of reference images from the remaining Celeb-HQ dataset. Table 1 shows our quantitative evaluation across different methods. Overall, our method outperforms previous approaches on the majority of metrics, attaining either the top or second position across all evaluated criteria. Specifically, the highest SSIM and IDS scores, along with a notable PSNR ranking, demonstrate that *Stable-Hair* effectively preserves both the background and identity of the source image during hairstyle transfer. Additionally, the lowest FID score indicates that our method achieves superior visual fidelity and realism in the transferred hairstyles. Achieving the highest CLIP-I score indicates that our method excels at aligning the transferred hairstyle with the reference image, showing strong performance in maintaining the desired style.

User Study. Considering the subjective nature of the hairstyle transfer task, we conducted a comprehensive user study involving 30 volunteers. Specifically, we randomly sampled 20 sets of data from our quantitative experiments

Method	CLIP-I \uparrow	FID \downarrow	PSNR \uparrow	SSIM \uparrow	IDS \uparrow
Ours (latent)	0.434	35.128	30.980	0.680	0.779
Ours (pixel)	0.423	39.236	29.317	0.668	0.770

Table 2: **Quantitative ablation** of different design options. ('pixel' refers to the use of the original ControlNet, while 'latent' denotes the use of the Latent ControlNet.)

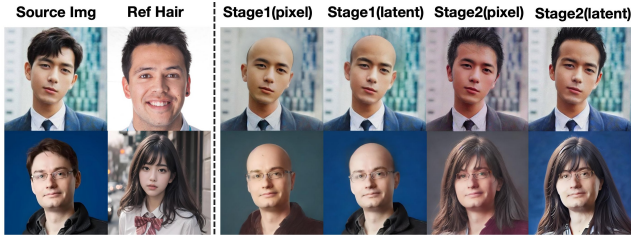


Figure 6: Visual ablation study results comparing different design options. The Latent ControlNet clearly maintains color consistency, whereas the original ControlNet introduces color discrepancies at each stage.

and selected 10 popular hairstyles from social media as reference styles, using a corresponding number of source images randomly sampled from FFHQ datasets to create an additional 10 sets of data using various algorithms. This resulted in a total of 30 triplets, each consisting of an original image, a reference image, and the transfer results. As with previous methods (Wei et al. 2022), the test results from different algorithms were randomized. For each test sample, volunteers were asked to choose the best option based on three criteria: transfer accuracy, preservation of unrelated attributes, and visual naturalness. The results shown in Table 3 indicate that our method outperforms the comparison methods in transfer accuracy, preservation of unrelated attributes, and visual naturalness.

Ablation Study. To thoroughly investigate the role of each module in our method, we conducted systematic ablation. As illustrated in Fig. 6, it is evident that models trained with pixel-conditioned input using ControlNet often exhibit color discrepancies (as seen in the first and third columns of the results), leading to inconsistencies between the source and target images. Using our proposed Latent ControlNet, which maps the hair removal and transfer process from pixel space to latent space, we effectively eliminate these color inconsistencies, significantly enhancing content preservation. The results of the quantitative ablation are shown in Table 2. Our method consistently outperforms across all metrics, highlighting the superior transfer capabilities and the enhanced preservation of identity and background achieved by the Latent ControlNet.

To further investigate the impact of video data on our method, we conducted training using only image data while keeping the same parameter settings. The results of the visual ablation study, as shown in Fig. 7, clearly indicate that when trained solely with image data, our method struggles to

Metrics	Barber	Hair	SYH	Hair	Hair	Ours
	shop	FastGAN		clip	clipV2	
Accuracy(%)	18.6	20.1	19.2	5.1	10.2	26.8
Preservation(%)	11.1	17.3	20.0	7.4	21.8	22.4
Naturalness(%)	13.4	20.2	15.9	11.2	18.7	20.6

Table 3: **User study** on hair transfer. Accuracy denotes the accuracy for hair transfer, Preservation indicates the ability to preserve irrelevant regions and Naturalness denotes the visual realism of the generated image. Our method achieves the best results in all three categories.




Figure 7: **Visual ablation** results of different training data. When trained exclusively on image data, the model often resorts to a simplistic copy-paste of the hairstyle. However, by integrating video data into the training process, our method gains pose robustness, enabling adaptive and natural hair transfer based on facial orientation.

handle significant variations in facial poses from the source images, leading to a collage-like effect with less natural transitions. In contrast, when trained with a mix of video data, our method demonstrates robustness to changes in facial pose. See the supplements for additional ablation studies.

Conclusions

In this paper, we introduce Stable-Hair, the first framework to tackle hairstyle transfer using diffusion techniques. This method marks a significant advancement, achieving stable and fine-grained real-world hairstyle transfers previously unattainable. Stable-Hair features a two-stage pipeline. The first stage uses a Bald Converter to transform the source image into a bald proxy image. Following this step, we designed an automated data collection pipeline that leverages the Bald Converter, an inpainting model, and ChatGPT to gather a diverse and large-scale dataset of triplets for training the second stage. In the second stage, a Hair Extractor and a Latent IdentityNet are utilized to accurately and robustly transfer the target hairstyle onto the bald image. Extensive experiments demonstrate that Stable-Hair achieves commercial-grade hairstyle transfer capabilities, setting a new standard in the field.

References

- Arar, M.; Gal, R.; Atzmon, Y.; Chechik, G.; Cohen-Or, D.; Shamir, A.; and Bermano, A. H. 2023. Domain-Agnostic Tuning-Encoder for Fast Personalization of Text-To-Image Models. *arXiv preprint arXiv:2307.06925*.
- Bodur, R.; Gundogdu, E.; Bhattarai, B.; Kim, T.-K.; Donoser, M.; and Bazzani, L. 2023. iEdit: Localised Text-guided Image Editing with Weak Supervision. *arXiv preprint arXiv:2305.05947*.
- Cao, M.; Wang, X.; Qi, Z.; Shan, Y.; Qie, X.; and Zheng, Y. 2023. MasaCtrl: Tuning-Free Mutual Self-Attention Control for Consistent Image Synthesis and Editing. *arXiv preprint arXiv:2304.08465*.
- Chang, S.; Kim, G.; and Kim, H. 2023. Hairerf: Geometry-aware image synthesis for hairstyle transfer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2448–2458.
- Chefer, H.; Alaluf, Y.; Vinker, Y.; Wolf, L.; and Cohen-Or, D. 2023. Attend-and-Excite: Attention-Based Semantic Guidance for Text-to-Image Diffusion Models. *arXiv:2301.13826*.
- Chung, C.; Kim, T.; Nam, H.; Choi, S.; Gu, G.; Park, S.; and Choo, J. 2022. Hairfit: pose-invariant hairstyle transfer via flow-based hair alignment and semantic-region-aware inpainting. *arXiv preprint arXiv:2206.08585*.
- Deng, J.; Guo, J.; Niannan, X.; and Zafeiriou, S. 2019. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *CVPR*.
- Epstein, D.; Jabri, A.; Poole, B.; Efros, A. A.; and Holynski, A. 2023. Diffusion self-guidance for controllable image generation. *arXiv preprint arXiv:2306.00986*.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion.
- Guo, X.; Kan, M.; Chen, T.; and Shan, S. 2022. GAN with Multivariate Disentangling for Controllable Hair Editing. *Springer, Cham*.
- 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. *Advances in neural information processing systems*, 30.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Jia, X.; Zhao, Y.; Chan, K. C.; Li, Y.; Zhang, H.; Gong, B.; Hou, T.; Wang, H.; and Su, Y.-C. 2023. Taming encoder for zero fine-tuning image customization with text-to-image diffusion models. *arXiv preprint arXiv:2304.02642*.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Khwanmuang, S.; Phongthawee, P.; Sangkloy, P.; and Suwajanakorn, S. 2023. StyleGAN Salon: Multi-View Latent Optimization for Pose-Invariant Hairstyle Transfer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kim, J.; Gu, G.; Park, M.; Park, S.; and Choo, J. 2023. StableVITON: Learning Semantic Correspondence with Latent Diffusion Model for Virtual Try-On. *arXiv preprint arXiv:2312.01725*.
- 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. *arXiv preprint arXiv:2208.07765*.
- Li, P.; Huang, Q.; Ding, Y.; and Li, Z. 2023. LayerDiffusion: Layered Controlled Image Editing with Diffusion Models. *arXiv preprint arXiv:2305.18676*.
- Li, S.; Zeng, B.; Feng, Y.; Gao, S.; Liu, X.; Liu, J.; Li, L.; Tang, X.; Hu, Y.; Liu, J.; et al. 2024. Zone: Zero-shot instruction-guided local editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6254–6263.
- Ma, W.-D. K.; Lewis, J.; Kleijn, W. B.; and Leung, T. 2023. Directed diffusion: Direct control of object placement through attention guidance. *arXiv preprint arXiv:2302.13153*.
- Mou, C.; Wang, X.; Song, J.; Shan, Y.; and Zhang, J. 2023a. DragonDiffusion: Enabling Drag-style Manipulation on Diffusion Models. *arXiv preprint arXiv:2307.02421*.
- Mou, C.; Wang, X.; Xie, L.; Zhang, J.; Qi, Z.; Shan, Y.; and Qie, X. 2023b. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*.
- Nikolaev, M.; Kuznetsov, M.; Vetrov, D.; and Alanov, A. 2024. HairFastGAN: Realistic and Robust Hair Transfer with a Fast Encoder-Based Approach. *arXiv preprint arXiv:2404.01094*.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. *arXiv:2103.00020*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. *arXiv:2204.06125*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2022. DreamBooth: Fine Tuning Text-to-image Diffusion Models for Subject-Driven Generation.
- Ruiz, N.; Li, Y.; Jampani, V.; Wei, W.; Hou, T.; Pritch, Y.; Wadhwa, N.; Rubinstein, M.; and Aberman, K. 2023. HyperDreamBooth: HyperNetworks for Fast Personalization of Text-to-Image Models. *arXiv preprint arXiv:2307.06949*.

- Saha, Rohit; Duke, Brendan; Shkurti, Florian; Taylor, Graham; Aarabi, and Parham. 2021. LOHO: Latent Optimization of Hairstyles via Orthogonalization. In *CVPR*.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E.; Ghasemipour, S. K. S.; Ayan, B. K.; Mahdavi, S. S.; Lopes, R. G.; Salimans, T.; Ho, J.; Fleet, D. J.; and Norouzi, M. 2022a. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *arXiv:2205.11487*.
- 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. 2022b. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35: 36479–36494.
- Shu, C.; Wu, H.; Zhou, H.; Liu, J.; Hong, Z.; Ding, C.; Han, J.; Liu, J.; Ding, E.; and Wang, J. 2022. Few-shot head swapping in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10789–10798.
- 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. *ACM Transactions on Graphics (TOG)*, 39(4): 1–13.
- Tsaban, L.; and Passos, A. 2023. Ledits: Real image editing with ddpm inversion and semantic guidance. *arXiv preprint arXiv:2307.00522*.
- Wang, R.; Guo, H.; Liu, J.; Li, H.; Zhao, H.; Tang, X.; Hu, Y.; Tang, H.; and Li, P. 2024. StableGarment: Garment-Centric Generation via Stable Diffusion. *arXiv preprint arXiv:2403.10783*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- 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. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Wei, T.; Chen, D.; Zhou, W.; Liao, J.; Zhang, W.; Hua, G.; and Yu, N. 2023. HairCLIPv2: Unifying Hair Editing via Proxy Feature Blending. *ICCV*.
- Wu, Y.; Yang, Y.-L.; and Jin, X. 2022. HairMapper: Removing Hair From Portraits Using GANs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4227–4236.
- Xie, D.; Wang, R.; Ma, J.; Chen, C.; Lu, H.; Yang, D.; Shi, F.; and Lin, X. 2023. Edit everything: A text-guided generative system for images editing. *arXiv preprint arXiv:2304.14006*.
- Xu, Y.; Gu, T.; Chen, W.; and Chen, C. 2024. Ootdiffusion: Outfitting fusion based latent diffusion for controllable virtual try-on. *arXiv preprint arXiv:2403.01779*.
- Yang, L.; Zeng, B.; Liu, J.; Li, H.; Xu, M.; Zhang, W.; and Yan, S. 2024. EditWorld: Simulating World Dynamics for Instruction-Following Image Editing. *arXiv preprint arXiv:2405.14785*.
- Zeng, J.; Song, D.; Nie, W.; Tian, H.; Wang, T.; and Liu, A. 2023. CAT-DM: Controllable Accelerated Virtual Try-on with Diffusion Model. *arXiv preprint arXiv:2311.18405*.
- Zhang, L.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*.
- Zhang, M.; and Zheng, Y. 2018. Hair-GANs: Recovering 3D Hair Structure from a Single Image. *arXiv e-prints*, arXiv:1811.06229.
- Zhang, S.; Huang, L.; Chen, X.; Zhang, Y.; Wu, Z.-F.; Feng, Y.; Wang, W.; Shen, Y.; Liu, Y.; and Luo, P. 2024a. FlashFace: Human Image Personalization with High-fidelity Identity Preservation. *arXiv preprint arXiv:2403.17008*.
- Zhang, Y.; Song, Y.; Liu, J.; Wang, R.; Yu, J.; Tang, H.; Li, H.; Tang, X.; Hu, Y.; Pan, H.; et al. 2024b. Ssr-encoder: Encoding selective subject representation for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8069–8078.
- Zhang, Y.; Song, Y.; Yu, J.; Pan, H.; and Jing, Z. 2024c. Fast Personalized Text to Image Synthesis with Attention Injection. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6195–6199. IEEE.
- Zhang, Y.; Wei, L.; Zhang, Q.; Song, Y.; Liu, J.; Li, H.; Tang, X.; Hu, Y.; and Zhao, H. 2024d. Stable-Makeup: When Real-World Makeup Transfer Meets Diffusion Model. *arXiv preprint arXiv:2403.07764*.
- Zhang, Z.; Han, L.; Ghosh, A.; Metaxas, D. N.; and Ren, J. 2023. Sine: Single image editing with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6027–6037.
- Zhao, S.; Chen, D.; Chen, Y.-C.; Bao, J.; Hao, S.; Yuan, L.; and Wong, K.-Y. K. 2023. Uni-ControlNet: All-in-One Control to Text-to-Image Diffusion Models. *arXiv preprint arXiv:2305.16322*.
- Zhu, P.; Abdal, R.; Femiani, J.; and Wonka, P. 2021. Barber-shop: GAN-based Image Compositing using Segmentation Masks. *arXiv:2106.01505*.
- Zhu, P.; Abdal, R.; Femiani, J.; and Wonka, P. 2022. HairNet: Hairstyle Transfer with Pose Changes.