

## Zero-Shot Head Swapping in Real-World Scenarios

Taewoong Kang<sup>1\*</sup> Sohyun Jeong<sup>1,2\*</sup> Hyojin Jang<sup>1\*</sup> Jaegul Choo<sup>1,2</sup>  
<sup>1</sup>KAIST <sup>2</sup>FLIPTION  
 {keh0t0, jsh0212, wkdgywlsrud, jchoo}@kaist.ac.kr



Figure 1. **Head-swapped images generated by our approach.** Using the proposed method, HID, the head in the images of *Head* column is seamlessly integrated onto the images of *Body* column, resulting in realistic and cohesive head-swapped images in the *Swapped* column.

### Abstract

With growing demand in media and social networks for personalized images, the need for advanced head-swapping techniques—integrating an entire head from the head image with the body from the body image—has increased. However, traditional head-swapping methods heavily rely

on face-centered cropped data with primarily frontal-facing views, which limits their effectiveness in real-world applications. Additionally, their masking methods, designed to indicate regions requiring editing, are optimized for these types of dataset but struggle to achieve seamless blending in complex situations, such as when the original data includes features like long hair extending beyond the masked area. To overcome these limitations and enhance adaptability in diverse and complex scenarios, we propose a novel head

swapping method, **HID**, that is robust to images including the full head and the upper body, and handles from frontal to side views, while automatically generating context-aware masks. For automatic mask generation, we introduce the **IOMask**, which enables seamless blending of the head and body, effectively addressing integration challenges. We further introduce the hair injection module to capture hair details with greater precision. Our experiments demonstrate that the proposed approach achieves state-of-the-art performance in head swapping, providing visually consistent and realistic results across a wide range of challenging conditions.

## 1. Introduction

With recent advancements in media and social networks, the demand for face-editing has increased, enabling users to replace faces with those of celebrities or fictional characters. To this end, continuous developments have been made in Face Swapping [8, 22, 29, 32, 49]. However, face swapping only replaces the face identity (face ID) that included eyes, nose, lip, eyebrow and skin, which limits its realism in applications like virtual avatar generation, movie and advertisement synthesis, and social media content creation, where differences in facial shape and hairstyle can make the results less convincing. In these cases, head swapping becomes essential.

Unlike face swapping, head swapping combines the entire head, including the face ID, face shape, and hairstyle, of the head image with the body image’s body, as their differences are highlighted in Fig. 2. Therefore, while face swapping only applies the face ID to the body image, head swapping needs to seamlessly merge the entire head from the head image with the body in the body image. Due to its increased complexity, head swapping still faces numerous unresolved challenges. 1) For a realistic head swap, it requires seamless blending between the swapped head and the original body. 2) Unlike face swapping, where only the face ID features of the head image are extracted, head swapping has to also handle the structural information of the head that includes a hairstyle and a facial structure.

These challenges stem from the use of face-centered cropped datasets with primarily frontal-facing views and masking methods. As shown in Fig. 3(a), most head-swapping methods [2, 14, 32] perform head swapping only on face-centered cropped data, which may not fully contain the hair. Additionally, the masking methods proposed in [2, 14, 32] are optimized for these types of images and, thus lacking the ability to create a seamless transition between the head and body. Therefore, if users require head swapping in images that include the full head as well as the body, previous methods need an additional phase to paste the head back into the body, often resulting in inharmonic

	Head Swap	Face Swap
Source	Face (ID and shape), Hairstyle, Skin	Face (ID)
Target	Pose, Background	Pose, Hairstyle, Face (Shape), Background

Figure 2. Face swapping simply applies the head image’s face ID to the body image. In contrast, head swapping requires applying not only the head image’s face ID but also the hairstyle, face shape, and skin tone.



(a) Examples of datasets used in the previous head swapping methods



(b) Examples of the dataset used in our approach

Figure 3. Most previous head-swapping methods [2, 14, 32] rely on datasets that require involving the zoomed-in, face-centered images as shown in (a). In this case, an additional step is needed to merge the head-swapped images back onto the original body, which can often result in mismatched or unnatural outcomes. In contrast, our approach uses the dataset that includes the whole upper body, as shown in (b), enabling head swapping in more realistic, real-world scenarios.

nous images. As this issue is clearly shown in Fig. 4, if the original head in the body images has long hair beyond the cropped area, remnants of the hair remain visible. Furthermore, datasets with primarily frontal-facing views lack robustness to diverse facial orientations. Although there have been attempts to address this issue, such as HeSer, it operates in a few-shot manner [40] by utilizing a wide range of view images extracted from video. This highlights the need for a more challenging dataset that includes upper body images with diverse facial orientations and has only a single image per person. To the best of our knowledge, no existing method efficiently addresses these real-world complexities in a zero-shot manner.

Therefore, we design a novel zero-shot head swapping approach, named Head Injection Diffusion (HID), with proposing Inverse-Orthogonal Mask (IOMask) and hair injection module, which are components of our approach to

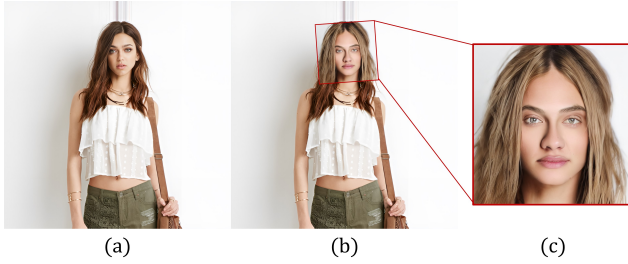


Figure 4. When existing methods that use face-centered cropped datasets [2, 14, 32] handle cases where the head-swapped image need to be pasted back with the original body image, they could generate inharmonious outcomes because the remaining parts in the body image cannot be addressed. The original body image (a) has brown hair, while the head-swapped image (c) has blonde hair, causing the final image (b) to show a noticeable color inconsistency in the hair.

perform head swapping in real-world scenarios with this challenging dataset. Our approach not only handles diverse facial orientations in a zero-shot setting but also leverages a SHHQ [12] dataset that showcases the entire upper body. This approach combines the strengths of existing datasets while addressing the limitations of previous head swapping methods as illustrated in Fig. 3. To tackle the remaining challenges in ensuring that the head of the head image blends naturally onto the body of the body image, we propose our IOMask, which automatically generates a context-aware mask. In head swapping tasks, automating the mask generation process is essential for real-world applications. Aligning the orientation of the head of the head image with head of the body image is crucial, but it is challenging for users to predict and provide such a mask manually and existing automatic masking methods have limitations as aforementioned. We address this issue by introducing the IOMask, extracted within HID, is obtained by leveraging the orthogonal component between the reconstruction condition and the editing condition, both of which originate from the inverted latent. Furthermore, to preserve face ID, face shape and hairstyle, we build our approach PhotoMaker V2 [1]<sup>1</sup>, a state-of-the-art identity-oriented personalization model. Inspired by this method, we introduce our hair injection module to transfer the hair information more precisely.

Our main contributions are as follows:

- We propose a zero-shot head swapping approach, HID, with introducing IOMask and hair injection module, which ensures seamless blending and better preserving intricate hairstyle details, leading to robust performance in real-world applications by using challenging dataset.
- We introduce IOMask, a novel approach that automatically generates realistic, context-aware masks for seam-

less head-body integration. This approach effectively blends the head with the body and removes any remains of the original body image, enhancing the practicality of head swapping for real-world applications.

- Through experimental results, we demonstrated state-of-the-art head swapping performance in handling complex scenarios.

## 2. Related work

### 2.1. Head Swap

Although there is relatively sparse research on head-swapping tasks, existing studies can be divided into two categories: few-shot approaches [32, 40] and zero-shot approaches [2, 14, 43]. Few-shot methods, such as those proposed by [32, 40], utilize videos to give a few shots of head input from different views so that the head in the head image is better aligned with the head in the body image. On the other hand, HSDiffusion [43] is a zero-shot diffusion-based method but just aligns the center points of the head and body images and cutting and pasting the components accordingly. In other words, it does not align the direction of the head in the head image with the direction of the head in the body image. Therefore, HSDiffusion does not precisely align with traditional head swapping task. Other zero-shot approaches, like FaceX [14] and REFace [2], also leverage diffusion models for head swapping. [14] is a unified method for handling diverse facial tasks and [2] is a face-swapping method adapted to support head swapping.

However, most of these existing head swapping methods [2, 14, 32] has a critical weakness in that they rely on face-centered cropped datasets with primarily frontal-facing view. Since head swapping is type of an inpainting task, generating an appropriate mask for the edited region is essential, but the mask generation process in [2, 14, 32] are optimized to these kinds of datasets. In practical applications, this often results in disharmonious outputs, especially when uncropped areas require adjustments for seamless head integration. Furthermore, even though HeSer [40] is not exactly corresponding to this problem, it is a few-shot based approach, which is less efficient than zero-shot. Our proposed approach performs in a zero-shot manner and uses dataset that contains upper body images with diverse facial orientations, enabling more practical and seamless head swapping in real-world applications.

### 2.2. Diffusion Models

Diffusion models [19, 37, 42] have achieved impressive advancements in generating images based on text prompts [21, 35, 36, 38] drawing significant attention in recent years. Their exceptional performance is largely due to the availability of high-quality, large-scale text-image datasets [5, 39], continuous improvements in foundational models [7,

<sup>1</sup>The code has been released but the paper has not been published yet.

31], enhancements in conditioning encoders, and better control mechanisms [3, 24, 28, 46, 47].

**Additional condition.** ControlNet [47] and T2I-Adapter [28] improve diffusion models by encoding spatial information, such as edges, depth, and human pose, to offer more control over the generated content through spatial conditioning alongside text prompts. ControlNet [47] achieves this by cloning the base model’s encoder and incorporating conditioning data as residuals into the model’s hidden states, enhancing coherence with spatial cues. IP-Adapter [46] further extends control by using high-level semantics from a reference image, projecting its embedding into the text encoder’s space to enable image generation influenced by both visual and textual prompts.

**Identity-oriented Diffusion Model.** Identity-oriented diffusion models [1, 25, 30, 44] enable efficient, high-fidelity identity synthesis. PhotoMaker [25] achieves this by fusing ID embedding with specific text embeddings. PhotoMaker V2 [1] improves upon PhotoMaker by utilizing more data and introducing an improved ID extraction encoder. InstantID [44] employs a cross-attention-based injection technique similar to IP-Adapter [46], applying it independently to text embeddings. Arc2Face [30] integrates face recognition features from ArcFace [10], fusing them before passing through the text encoder.

### 2.3. Image Editing

Recently, leveraging the strong generative capabilities of diffusion models, several studies have aimed to facilitate image editing that aligns closely with user intent [3, 9, 16, 21, 27]. One critical aspect of image editing is local editing, where specific areas of an image are modified based on user guidance. Many methods [9, 15, 27] address this by employing masks to define the regions requiring adjustment.

In general, masks for local editing within diffusion models can be obtained through two main approaches: one using difference between predicted noises and the other leveraging attention-based [6, 16] techniques. Using predicted noise’s difference methods [9, 15, 27] use of the disagreement in predictions of stable diffusion [37] with source and target captions. This method is advantageous due to their computational efficiency, providing faster editing capabilities. However, they can suffer from inaccuracies, which we aim to mitigate with our proposed approach. On the other hand, attention-based techniques [4, 6, 13, 16] are typically more precise but have the drawback of slower processing speeds. Although segmentation masks could also be used to delineate editing areas, they come with limitations in certain transformations, such as changing short hair to long hair, where they may lack adaptability. Our method

addresses these challenges by improving the accuracy, providing a balanced solution that enhances both precision and speed in local editing tasks.

## 3. Method

We propose a novel approach, HID, that takes a body image  $I_b$  and a head image  $I_h$  as inputs and outputs a swapped image  $I_o$ , where the head in  $I_b$  is seamlessly replaced with the one in  $I_h$ . We also introduce IOMask for precise mask extraction to guide head integration (Sec.3.2) and utilizes hair injection module to generate heads that match the input head images (Sec.3.3). Together, these components form our cohesive head-swap method (Sec.3.4). Fig. 5 provides an overview of our process.

### 3.1. Preliminary

**Classifier-free guidance.** In text-guided generation, a significant challenge is managing the enhanced influence of the conditioned text on the output. To address this, [18] introduced a technique known as classifier-free guidance, CFG, which combines an unconditional prediction with a conditioned prediction to achieve the final result. Formally, let  $\emptyset = \psi(“”)$  represent the embedding of an empty or null text, and let  $w$  denote the guidance scale parameter. The CFG prediction is then formulated as follows:

$$\epsilon_{\theta}(z_t, t, C, \emptyset) = \epsilon_{\theta}(z_t, t, \emptyset) + w \cdot (\epsilon_{\theta}(z_t, t, C) - \epsilon_{\theta}(z_t, t, \emptyset)). \quad (1)$$

**DDIM inversion.** A simple inversion method has been suggested for DDIM sampling [41], which relies on the premise that the ordinary differential equation process can be retraced in reverse if the step size is sufficiently small.

$$z_{t+1} = \sqrt{\frac{\alpha_{t+1}}{\alpha_t}} z_t + \left( \sqrt{\frac{1}{\alpha_{t+1}} - 1} - \sqrt{\frac{1}{\alpha_t} - 1} \right) \cdot \epsilon_{\theta}(z_t, t, C). \quad (2)$$

Put differently, the diffusion process follows a reversed trajectory, going from  $z_0$  to  $z_t$  rather than from  $z_t$  to  $z_0$ , with  $z_0$  initialized as the encoding of a real input image. Moreover, DDIM Inversion uses CFG with  $w = 1$  that means text-conditioned predicted noise  $\epsilon_{\theta}(z_t, t, C)$  is used.

**PhotoMaker V2.** Given a few ID images to be customized, PhotoMaker V2 [1] generates a new photo-realistic human image that retains the characteristics of the input IDs. This capability is achieved by updating text embeddings. Specifically, the text embedding for a particular class is fused with the image embeddings extracted by CLIP [34] image encoder and the ID embeddings extracted by ID encoder [20]. In this paper, we refer the model that creates the fused embedding, as ID Fusion model. The output of ID Fusion model replaces the class text embedding before it is input into the U-Net.

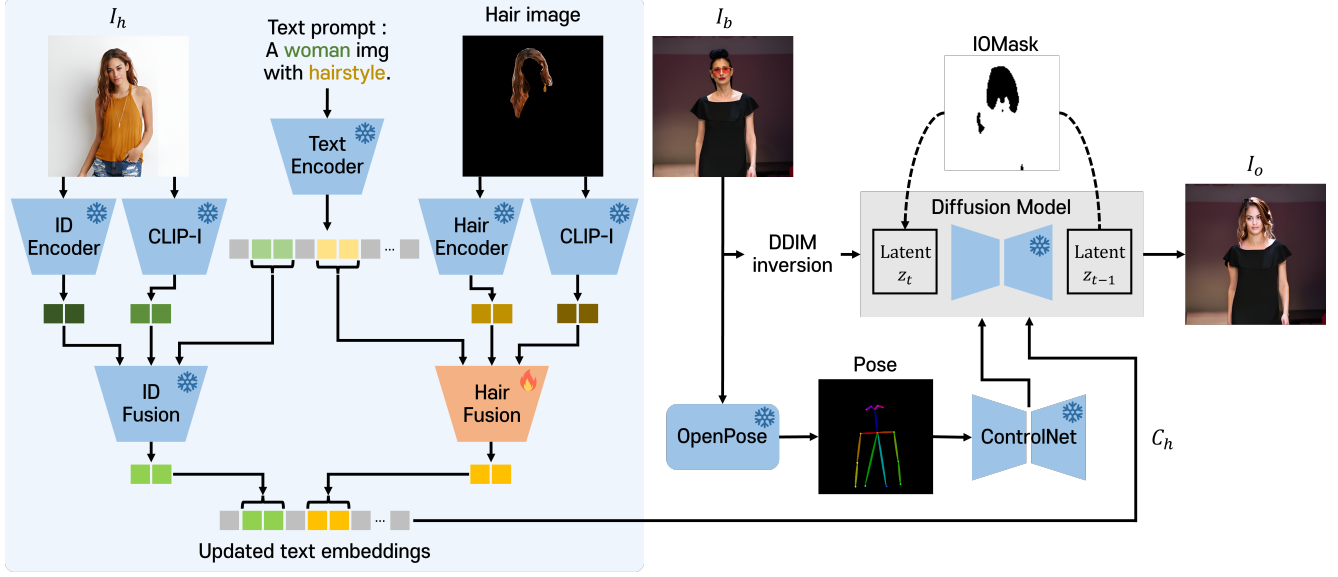


Figure 5. **Overview of HID.** Our HID consists of two main stages (left and right). In the left stage (blue region), we obtain updated text embeddings by fusing embeddings. These fused ID embeddings and fused hair embeddings replace the part of original text embeddings, resulting in updated text embeddings. In the right stage (white region), the final output  $I_o$ , a head swapped image, is generated. DDIM inversion is performed to reconstruct the image while leveraging our IOMask to infer which parts of the body image should be removed, thereby generating the head image. During this process, the updated text embeddings obtained in the left stage, along with the output from ControlNet, serves as conditioning for the diffusion model.

### 3.2. IOMask

For accurate head-swapping, obtaining a mask for head region is crucial to generate on the body image  $I_b$ . Inspired by [9, 15, 27], we achieve this by extracting the mask within our model. Given an body image  $I_b$ , we first apply DDIM [41] Inversion with our model, obtaining the latent  $\hat{z}_t$  at a specific time step  $t$ . If we use DDIM Inversion with guidance scale  $w = 1$ , that provides a faithful reconstruction when adapting the denoising guidance keep the scale same. For specific time-step  $t$ , we can get  $\epsilon_\theta(\hat{z}_t, t, C_b)$  from denoising  $\hat{z}_t$ , where  $C_b$  means body image’s text embeddings, and established this for a baseline. The Inverse-Orthogonal map, called IO map, is then determined by the areas that differ under head-conditioned predicted noise  $\epsilon_\theta(\hat{z}_t, t, C_h, \emptyset)$ , where  $C_h$  means head image’s text embeddings with fused ID embeddings and fused hair embeddings. Note that CFG is conducted for head-conditioned predicted noise. For simplicity, we’ll omit notation  $\hat{z}_t$  and  $t$ .

To obtain IO map, we compare  $\epsilon_\theta(C_h, \emptyset)$  and  $\epsilon_\theta(C_b)$ . Using  $\epsilon_\theta(C_b)$  as the reference, we compute the mask by isolating vector components of  $\epsilon_\theta(C_h, \emptyset)$  orthogonal to the body condition, rather than relying on absolute differences  $\epsilon_\theta(C_h, \emptyset) - \epsilon_\theta(C_b)$ . We provide a visualization of this process in vector form in Fig. 6. When adjusting normalization and thresholding, the distribution of value magnitudes is critical. Values aligned in the similar direction should be smaller, while those in differing directions should be larger,

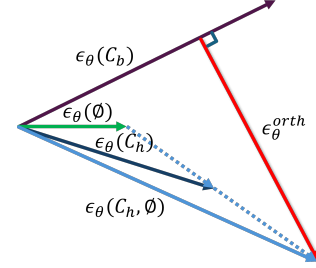


Figure 6. Visualization of obtaining  $\epsilon_\theta^{orth}$  in vector form.

as the head area will be oriented differently and should be emphasized. Therefore, we base our calculations on the components orthogonal to the reference  $\epsilon_\theta(C_b)$ . Thus, the formulation is given by:

$$\epsilon_\theta^{orth} = \epsilon_\theta(C_h, \emptyset) - \frac{\langle \epsilon_\theta(C_b), \epsilon_\theta(C_h, \emptyset) \rangle}{\|\epsilon_\theta(C_b)\|^2} \epsilon_\theta(C_b) \quad (3)$$

After normalizing the IO map between 0 and 1, we observe the presence of dotted artifacts. To address this issue and improve robustness, we apply a Gaussian filter. The Inverse-Orthogonal Mask, called IOMask,  $\mathcal{M}$  is then computed as follows:

$$\mathcal{M} = \mathbf{1}_{(GF(|\epsilon_\theta^{orth}|) \geq \tau)}, \quad (4)$$

where  $GF$  denotes the Gaussian filter function, and  $\tau \in [0, 1]$  is the threshold parameter.

### 3.3. Hair Injection Module

While PhotoMaker V2 [1] can generate human images that maintain a given identity to varying degrees, it cannot be directly applied to the head-swapping task because it is fundamentally an image generation model, not an image editing model. To address this, we first introduce the IOMask to provide the model with information about the region that needs to be edited.

A remaining challenge is to generate hairstyles that accurately reflect the hairstyle features in  $I_h$ . Since PhotoMaker V2 relies solely on text prompts for conditioning, aside from ID, we inject hairstyle information from images specifically for the head-swapping task by training the Hair Fusion model inspired by PhotoMaker V2. We train the Hair Fusion model with text prompts that includes 'A man/woman *img* with *hairstyle*.' *img* serves as a trigger word for ID injection, while the embeddings of *hairstyle* is fused with the hair embeddings extracted from the pre-trained Hair Encoder [50] and the hair image embeddings extracted from the CLIP image encoder. The Hair Fusion model comprises a Q-former and MLPs, and the text embeddings are updated such that the fused ID embeddings and fused hair embeddings replace the corresponding original text embeddings. This allows the model to learn a distribution that enables these updated embeddings to condition the UNet, facilitating effective head-swapping. Also, during hair injection module training, SCHP [23] was used to extract a mask for regions excluding the background, and a masked loss was applied to reconstruct the person, allowing for effective hairstyle integration.

### 3.4. Head Injection Diffusion Model

The Head Injection Diffusion model (HID) performs head swapping by combining the hair injection module with IOMask. Given a body image  $I_b$  and a head image  $I_h$ , the method outputs a swapped image  $I_o$ . To achieve this,  $I_h$  is processed through the ID injection module and a hair image is processed through hair injection module. Each module outputs a fused embedding which replaces the corresponding text embeddings to generate the head condition  $C_h$ , used as input through cross-attention. Additionally, the body pose is provided to the model via open-pose ControlNet [47] to ensure alignment.

The denoising process begins with the inversion latent  $\hat{z}_T$ , stored during IOMask generation. Starting with this latent ensures that the masked region does not begin without body details, thereby preventing mismatches in skin tones or clothing artifacts. This inversion latent is continuously blended with the mask during updates, promoting a more natural result. Each denoising stage refines a noisy latent  $z_t$  to obtain  $z_{t-1}$ .

During each denoising step, given the head condition  $C_h$ ,

the latent  $z_t$  is denoised to latent  $\tilde{z}_t$  and updated as follows:

$$z_{t-1} = \tilde{z}_{t-1} \odot \mathcal{M} + \hat{z}_{t-1} \odot (1 - \mathcal{M}) \quad (5)$$

By replacing the unmasked pixels with the inversion latent of the input image, we prevent the generation process from altering any pixels outside the mask. After iterative denoising, the head-swapped image  $I_o$  is obtained.

## 4. Experiments

### 4.1. Experimental Setup

**Dataset.** We use the SHHQ-1.0 dataset [12], consisting of 40K images, for training and evaluation. The original resolution of the dataset is  $1024 \times 512$ ; however, to focus on the head swapping task, we preprocess the images to  $512 \times 512$  by cropping out the lower part of the body. Given that the maximum resolution that SDXL [33] can generate is  $1024 \times 1024$ , we upsample the images to this resolution using GFPGAN v1.4 [45]. After image preprocessing, we obtain a total of 39,804 images, which are split in a 9:1 ratio. This results in 35,823 images for training and 3,981 images for evaluation. In addition, We generate image captions for training using the state-of-the-art image captioning model, LLaVa-NEXT [26], which is based on LLaMa-3-8B [11]. To incorporate ID and hairstyle information, we modify the captions to include the phrase 'A man/woman *img* with *hairstyle*.'

**Baselines.** To the best of our knowledge, there are no existing baselines that align precisely with our head swapping task, which operates in a zero-shot manner, reenacts the head in the head image to match the orientation of the head in the body image, and does not require face-centered cropping or alignment. Therefore, we compare our performance solely with REFace [2], the only head swapping method that has officially released its code between zero-shot head swapping method, to demonstrate that our approach outperforms an existing method.

**Evaluation.** All images in the evaluation dataset are used as head images, which are randomly paired with other images from the same dataset to serve as body images. We use 'A man/woman image with hairstyle' as a text prompt for the evaluation. The evaluation is conducted using the following metrics: Learned Perceptual Image Patch Similarity (LPIPS)[48], CLIP-I [34] and Frechet Inception Distance (FID)[17]. To demonstrate the performance of our approach, we calculate LPIPS and CLIP-I using images with masked hair or masked head regions, where the head mask includes both hair and face. We first generate human parsing masks using SCHP [23] and use these masks to isolate hair or head regions. After that, we measure LPIPS and CLIP-I between the hair and head regions of the generated images and the corresponding regions in the real im-

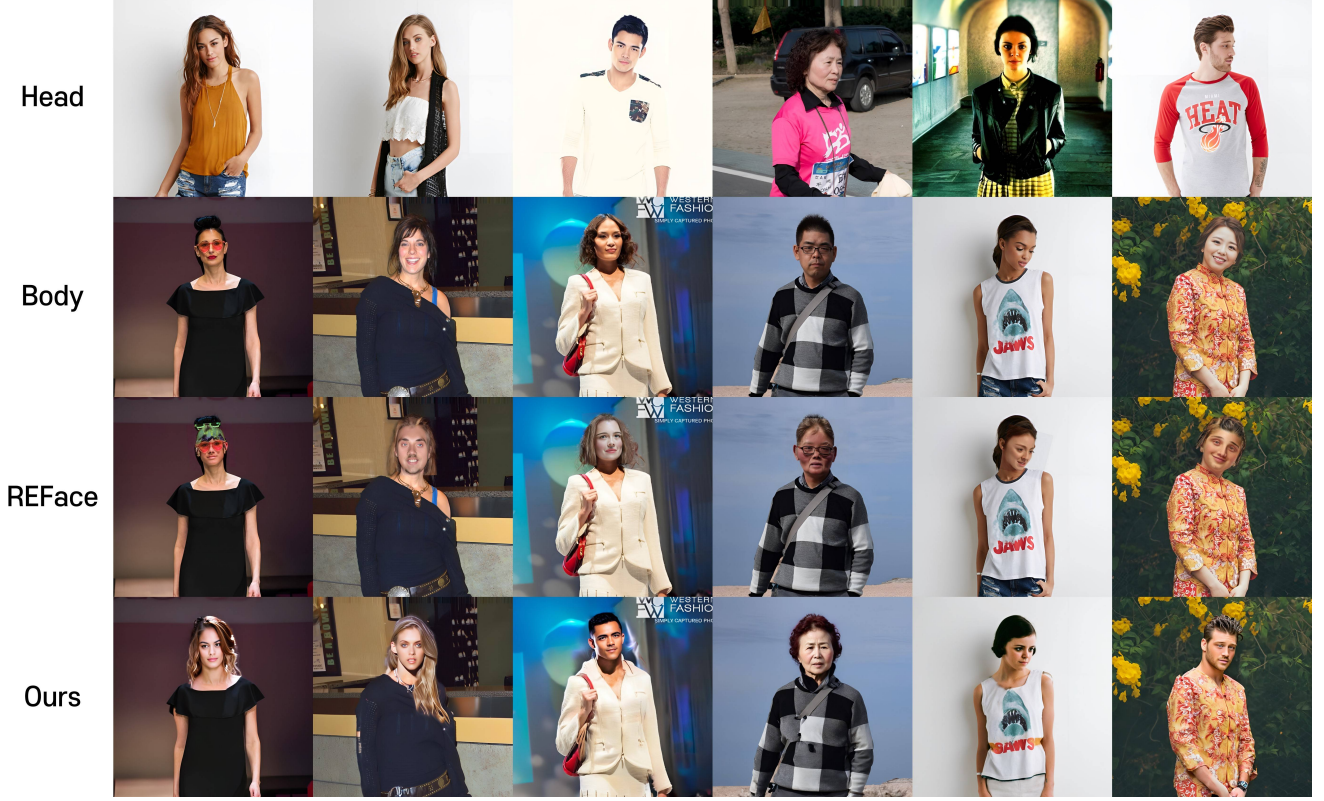


Figure 7. **Qualitative comparison.** The images in the *Head* row are combined with those in the *Body* row. The last two rows are the head-swapped results produced by each method.

ages from the evaluation dataset. This allows us to more closely compare the generated image with specific features of the real image to make an assessment of how similar and consistent the targeted areas are. Additionally, we use FID to assess the quality of generated images. However, since REFace [2] handles solely on face-centered cropped regions, a fair comparison is difficult. To address this, we crop the same region from our generated images to match the face-centered area used by REFace.

## 4.2. Results

**Quantitative results.** As shown in Tab. 1, our HID outperforms REFace [2] in both LPIPS and CLIP-I metrics for the head and hair regions. This result indicates that HID more accurately preserves ID and hairstyle in head images, which is essential for the head-swapping task. Moreover, we outperformed REFace on the FID score. This demonstrates that our generated images have superior quality, achieving a more natural integration of head and body features compared to existing methods.

**Qualitative results.** As shown in Fig. 2, the swapped head must preserve key attributes of the head image, including skin tone, hairstyle, identity, and face shape. Fig. 1 illustrates how our method effectively achieves these requirements, performing the head swap task while maintain-

ing these essential characteristics. Additionally, thanks to the IOMask, we can generate appropriate masks that allow the creation of long hair where needed, ensuring that only the relevant head region is modified while leaving other areas unchanged. In Fig. 7, we demonstrate that our method outperforms REFace [2], particularly in challenging cases where head orientation and gender differ significantly. Unlike REFace, which suffers from limitations inherent to the cropping approach—such as hair being cut off and not blending well with surrounding areas—our approach overcomes these issues, maintaining a seamless and coherent appearance even under these challenging conditions. Furthermore, unlike the mask used in REFace, our proposed IOMask effectively removes only the necessary face and hair areas from the body image, enabling a smooth and natural swap.

## 4.3. Ablation Study

### 4.3.1. IOMask

We demonstrate the effectiveness of our IOMask through an ablation study, comparing the results of three configurations: the naive IO map, IO map without orthogonal components, and the full IO map. The naive IO map is defined as  $\epsilon_\theta = \epsilon_\theta(C_h, \emptyset) - \epsilon_\theta(C_b, \emptyset)$ , where the ID embeddings of both head and body conditions are without further re-

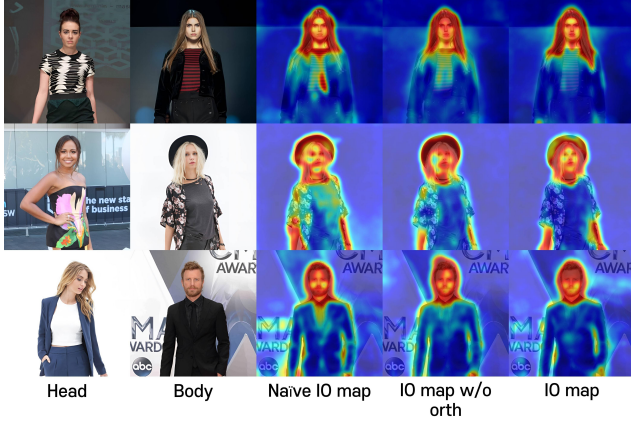


Figure 8. **Ablation study on IOMask configurations.** We compare three variants: the naive IO map, the IO map without orthogonal filtering, and the full IO map. The full IO map demonstrates improved focus and precision in identifying relevant areas for head swapping, as indicated by regions of high values in red.

finement. The IO map without orthogonal components is defined as  $\epsilon_{\theta} = \epsilon_{\theta}(C_h, \emptyset) - \epsilon_{\theta}(C_b)$ , omitting the orthogonal filtering step present in our full IO map. We visualized the IO map by overlaying it on body images, with red indicating regions of high values. As illustrated in Fig. 8, the naive IO map introduces substantial randomness and often covers irrelevant areas. In contrast, both the IO map without orthogonal components and the full IO map provide more relevant coverage. Notably, the full IO map focuses more precisely on the regions that have to be swapped.

#### 4.3.2. HID

We evaluate the performance impact of each element in the HID to show the effect of our methods as shown in Fig. 9. So, we sequentially remove the hair injection module and IO Mask, while applying both model with ControlNet and same input text, 'A man/woman img with hairstyle.'

Starting with the our full approach, which includes all components, we achieved optimal results with high-quality head swaps, including detailed hairstyle transfer and effective integration of the head with the body. When the hair injection module was removed, the generated results showed a noticeable reduction in hairstyle detail, as the method could no longer fully transfer the hairstyle from the head image onto the body image. This is because, while PhotoMaker V2 effectively maintains faceID, it does not inherently ensure hairstyle preservation.

Finally, although generation began from the inversion latent, removing the IO Mask resulted in retaining only a very slight amount of the body image and losing its ability to maintain regions of the body image outside the head. This setup led to inconsistencies in other areas of the body image, underscoring the importance of each component for achieving a natural and coherent final output.



Figure 9. **Ablation study on components of our approach.** We conducted ablation experiments by removing each component of our method individually. Without the hair injection module, the model fails to retain the head image’s hairstyle, as seen in the first row where the hair appears short. Additionally, removing the IO mask results in a complete loss of body image information.

	FID↓	Head		Hair	
		LPIPS↓	CLIP-I↑	LPIPS↓	CLIP-I↑
REFace	40.7162	0.0770	0.7867	0.0658	0.8563
Ours	<b>37.1879</b>	<b>0.0721</b>	<b>0.8512</b>	<b>0.0596</b>	<b>0.8686</b>

Table 1. **Quantitative comparison.** Ours outperforms REFace [2] across all metrics.

## 5. Conclusion

We propose the HID approach, an efficient head swapping method aimed at real-world applications. Our approach allows the head image to blend seamlessly with the body image while preserving original details, resulting in a natural outcome. Despite these significant improvements over previous approaches, there are still areas that require refinement. For example, regions beyond the head, like the hands, neck, or other skin areas, may experience unintended changes. In the Fig. 1, the second example from the right shows a case where the hands were masked to adjust their color, allowing them to match the skin tone of the head image. However, while the skin tone was transferred successfully, the pose and fine details of the original body image’s hands were not preserved, which highlights a limitation of our approach. Unlike previous methods that rely heavily on cropping or head-centric masking techniques, our approach takes into account which areas should be edited when conditioned with the head image, allowing for more flexibility and potential for improvement. Nevertheless, our method shows potential for further advancements, and future work could focus on better preserving details of these regions while still achieving seamless skin tone transfer.

## Acknowledgement

This work was supported by Institute for Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (RS-2019-II190075, Artificial Intelligence Graduate School Program(KAIST)). This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2025-00555621). We appreciate the high-performance GPU computing support of HPC-AI Open Infrastructure via GIST SCENT.

## References

- [1] Tencent ARC. PhotoMaker V2. <https://github.com/TencentARC/PhotoMaker>, 2024. Accessed: 2024-08-10. 3, 4, 6
- [2] Sanoojan Baliah, Qinliang Lin, Shengcai Liao, Xiaodan Liang, and Muhammad Haris Khan. Realistic and efficient face swapping: A unified approach with diffusion models. *arXiv preprint arXiv:2409.07269*, 2024. 2, 3, 6, 7, 8
- [3] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 4
- [4] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22560–22570, 2023. 4
- [5] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3558–3568, 2021. 3
- [6] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models, 2023. 4
- [7] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, et al. Pixart- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis. *arXiv preprint arXiv:2310.00426*, 2023. 3
- [8] Renwang Chen, Xuanhong Chen, Bingbing Ni, and Yanhao Ge. Simswap: An efficient framework for high fidelity face swapping. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2003–2011, 2020. 2
- [9] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022. 4, 5
- [10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. 4
- [11] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 6
- [12] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen-Change Loy, Wayne Wu, and Ziwei Liu. Stylegan-human: A data-centric odyssey of human generation. *arXiv preprint, arXiv:2204.11823*, 2022. 3, 6
- [13] Qin Guo and Tianwei Lin. Focus on your instruction: Fine-grained and multi-instruction image editing by attention modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6986–6996, 2024. 4
- [14] Yue Han, Jiangning Zhang, Junwei Zhu, Xiangtai Li, Yanhao Ge, Wei Li, Chengjie Wang, Yong Liu, Xiaoming Liu, and Ying Tai. A generalist facex via learning unified facial representation. *arXiv preprint arXiv:2401.00551*, 2023. 2, 3
- [15] Sai Sree Harsha, Ambareesh Revanur, Dhwanit Agarwal, and Shradha Agrawal. Genvideo: One-shot target-image and shape aware video editing using t2i diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7559–7568, 2024. 4, 5
- [16] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 4
- [17] 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 *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2017. 6
- [18] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 4
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [20] Deep Insight. InsightFace: 2D and 3D Face Analysis Project. <https://github.com/deepinsight/insightface>, 2023. Accessed: 2024-08-10. 4
- [21] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6007–6017, 2023. 3, 4
- [22] Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, and Fang Wen. Faceshifter: Towards high fidelity and occlusion aware face swapping. *arXiv preprint arXiv:1912.13457*, 2019. 2
- [23] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing, 2019. 6
- [24] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22511–22521, 2023. 4
- [25] Zhen Li, Mingdeng Cao, Xintao Wang, Zhongang Qi, Ming-Ming Cheng, and Ying Shan. Photomaker: Customizing re-

- alistic human photos via stacked id embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8640–8650, 2024. 4
- [26] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024. 6
- [27] Ashkan Mirzaei, Tristan Aumentado-Armstrong, Marcus A Brubaker, Jonathan Kelly, Alex Levinshtein, Konstantinos G Derpanis, and Igor Gilitschenski. Watch your steps: Local image and scene editing by text instructions. In *European Conference on Computer Vision*, pages 111–129. Springer, 2024. 4, 5
- [28] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4296–4304, 2024. 4
- [29] Yuval Nirkin, Yosi Keller, and Tal Hassner. Fsgan: Subject agnostic face swapping and reenactment. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7184–7193, 2019. 2
- [30] Foivos Paraperas Papantoniou, Alexandros Lattas, Stylianos Moschoglou, Jiankang Deng, Bernhard Kainz, and Stefanos Zafeiriou. Arc2face: A foundation model of human faces. *arXiv preprint arXiv:2403.11641*, 2024. 4
- [31] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 4
- [32] Ivan Perov, Daiheng Gao, Nikolay Chervoniy, Kunlin Liu, Sugasa Marangonda, Chris Umé, Dpfks, Carl Shift Facenheim, Luis RP, Jian Jiang, Sheng Zhang, Pingyu Wu, Bo Zhou, and Weiming Zhang. Deepfacelab: Integrated, flexible and extensible face-swapping framework, 2021. 2, 3
- [33] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 6
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 4, 6
- [35] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1 (2):3, 2022. 3
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3, 4
- [38] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 3
- [39] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 3
- [40] Changyong Shu, Hema Wu, Hang Zhou, Jiaming Liu, Zhibin Hong, Changxing Ding, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. Few-shot head swapping in the wild, 2022. 2, 3
- [41] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 4, 5
- [42] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [43] Qinghe Wang, Lijie Liu, Miao Hua, Pengfei Zhu, Wangmeng Zuo, Qinghua Hu, Huchuan Lu, and Bing Cao. Hs-diffusion: Semantic-mixing diffusion for head swapping, 2023. 3
- [44] Qixun Wang, Xu Bai, Haofan Wang, Zekui Qin, Anthony Chen, Huaxia Li, Xu Tang, and Yao Hu. Instantid: Zero-shot identity-preserving generation in seconds. *arXiv preprint arXiv:2401.07519*, 2024. 4
- [45] Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 6
- [46] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 4
- [47] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 4, 6
- [48] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [49] Yuhao Zhu, Qi Li, Jian Wang, Cheng-Zhong Xu, and Zhenan Sun. One shot face swapping on megapixels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4834–4844, 2021. 2
- [50] Adéla Šubrtová, Jan Čech, and Vojtěch Franc. Hairstyle transfer between face images. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pages 1–8, 2021. 6