

Few-Shot Head Swapping in the Wild

Changyong Shu¹ Hema Wu² Hang Zhou^{1*} Jiaming Liu^{1*} Zhibin Hong¹
Changxing Ding² Junyu Han¹ Jingtuo Liu¹ Errui Ding¹ Jingdong Wang¹

¹Department of Computer Vision Technology (VIS), Baidu Inc., ²South China University of Technology
{zhouhang09, liujiaming03, hongzhibin, hanjunyu, liujingtuo, dingerrui, wangjingdong}@baidu.com,
changyong.shu89@gmail.com, {201830252427@mail., chxding@}scut.edu.cn.



Figure 1. **Head swapping results generated by HeadSwapper.** The first line is the source image. The second line is the target image. The head swapping results are shown in the third line, where the source head is seamlessly transferred to the target in the wild.

Abstract

The head swapping task aims at flawlessly placing a source head onto a target body, which is of great importance to various entertainment scenarios. While face swapping has drawn much attention, the task of head swapping has rarely been explored, particularly under the few-shot setting. It is inherently challenging due to its unique needs in head modeling and background blending. In this paper, we present the **Head Swapper (HeSer)**, which achieves few-shot head swapping in the wild through two delicately designed modules. Firstly, a **Head2Head Aligner** is devised to holistically migrate pose and expression information from the target to the source head by examining multi-scale information. Secondly, to tackle the challenges of skin color variations and head-background mismatches in the swapping procedure, a **Head2Scene Blender** is introduced to simultaneously modify facial skin color and fill mismatched gaps on the background around the head. Particularly, seamless blending is achieved with the help of a **Semantic-Guided Color Reference Creation** procedure and a **Blending UNet**. Extensive experiments demonstrate that the pro-

posed method produces superior head swapping results on a variety of scenes.

1. Introduction

Human cognition of identity appearance is profoundly affected by not only facial structures but also head shapes and hairstyles. Head swapping, the ability to seamlessly replace the head in a target image with a source one (as shown in Fig. 1) would be of great importance to a variety of scenarios such as movie and advertisement composition, virtual humans creation, and deepfake video detection, etc.

While face swapping has long been a topic of interest [3, 17, 21, 23], only a few studies have been carried out on the task of head swapping. DeepFaceLab [23] requires large manual intervention to generate head swapping results, and they are totally incapable of handling mismatched regions. StylePoseGAN [24] tends to change the color of body skin and background in the target image in an undesired manner. Moreover, both methods fail to address few-shot head swapping, particularly for in-the-wild scenes.

We identify several properties that make few-shot head swapping more challenging than face swapping: 1) Head

*Corresponding authors.

swapping requires not only perfect facial identity and expression modeling, but also capturing the structural information of a whole head and the non-rigid hair. Thus previous identity extraction strategies for face swapping [3, 21] cannot be directly applied to head swapping. 2) There would be a huge *region mismatch between swapped head edges and backgrounds* caused by the editing of head shapes and hairstyles. Such a problem does not exist in the face swapping setting. 3) Moreover, similar to face swapping, the color difference between source and target skins needs to be handled carefully.

In this paper, we propose a framework called **Head Swapper (HeSer)**, which generates high fidelity head swapping results in the wild based on a few frames. Our key insight is to *positionally and emotionally align the source head with the target in a unified blender that seamlessly handles both color and background mismatches*. Two modules, namely the *Head2Head Aligner* and the *Head2Scene Blender* are devised. The *Head2Head Aligner* is responsible for finding a latent representation of a whole head as well as the facial details. It aligns the source head to the same pose and expression as the target image in a holistic manner. The identity, expression, and pose information are prominently balanced in a style-based generator by encoding multi-scale local and global information from both images. Moreover, a subject-specific fine-tuning procedure could further improve identity preservation and pose consistency.

To further blend the aligned head into the target scene, we devise a module named *Head2Scene Blender*, which provides both 1) color guidance for facial skins and 2) padding priors for inpainting the gaps on the background around the head. Thus both the skin color and edge-backgrounds mismatches between source and target can be handled within one unified sub-module. It efficiently creates colored references by building the correlations between pixels of the same semantic regions. Then with a blending UNet, seamless and realistic head swapping results can be rendered.

We summarize our main contributions as follows: **1)** We introduce a *Head2Head Aligner* that holistically migrates position and expression information from the target to the source head by examining multi-scale information. **2)** We design a *Head2Scene Blender* to simultaneously handle facial skin color and background texture mismatches. **3)** Our proposed **Head Swapper (HeSer)** produces photo-realistic head swapping results on different scenes. To the best of our knowledge, this is one of the earliest methods to achieve few-shot head swapping in the wild.

2. Related Work

The proposed two-stage Head Swapper is closely related to the topics of face swapping, head reenactment and image blending, which will be discussed below.

Face and Head Swapping. Many methods have been proposed for face swapping [1, 3, 18, 21, 22, 33, 34, 44]. While few studies tackle the task in a source-oriented manner [21, 22] that aims at reenacting and blending the source into the target, most methods are performed in a target-oriented manner [1, 3, 18, 34, 44]. They extract identity representation from the source image and inject it into the target image in a generative model. Face swapping approaches produce immutable hairstyles and head shapes, which limits the overall similarity between the generated results and the source.

The task of head swapping has been rarely studied. Deepfacelab [23] is the first work to tackle head swapping. However, it suffers from the following issues: 1) huge amounts of source data are required; 2) color transfer can only be poorly performed; and 3) the regions that require inpainting the fusion look unnatural. Later StylePoseGAN [24], which is specially designed for re-rendering using pose-/appearance-conditioned StyleGAN [29], potentially has the capability of head swapping. However, it tends to perform poorly for in-the-wild scenarios. Our proposed head swapping pipeline is the first work committed to achieving few-shot head swapping in the wild.

Few-Shot Head Reenactment. Head swapping requires placing the source head at the target head’s position with the same pose and expression, which is similar to the task of talking head generation and head reenactment [4, 6, 7, 14, 16, 19, 20, 26–28, 30, 38, 41, 43]. Here we will briefly discuss the visual-driven few-shot reenactment methods.

Recent few-shot methods can be roughly divided into warping-based and reconstruction-based methods. Warping-based approaches [26–28] deform the source images to imitate the motion of driving ones, they tend to work poorly for the case with a big range of motion. Reconstruction-based approaches mostly leverage generative adversarial networks (GANs) [8]. A great number of studies use intermediate representations such as landmarks [4, 36] and 3D models [41]. However, the inaccuracy of the structural information might lead to error accumulation. Certain recent studies focuses on the extraction and disentanglement of pose and identity information in the latent space [6, 19, 42, 43]. Particularly, LPD [6] verifies that a style-based generator can explicitly handles such information in latent vectors. Inspired by this type of work, we introduce the *Head2Head Aligner* to find the latent representation of the whole head as well as the facial details in a reconstruction-based learning pipeline.

Color Transfer. As we need to transfer the skin color of the source face to the target one, it is similar to the task of makeup transfer. Early makeup transfer methods utilize facial landmarks or parsing results as prior information [2, 9, 12, 13, 31], however, they cannot keep the facial expressions unchanged. The subsequently PSGAN [15, 25] and SCGAN

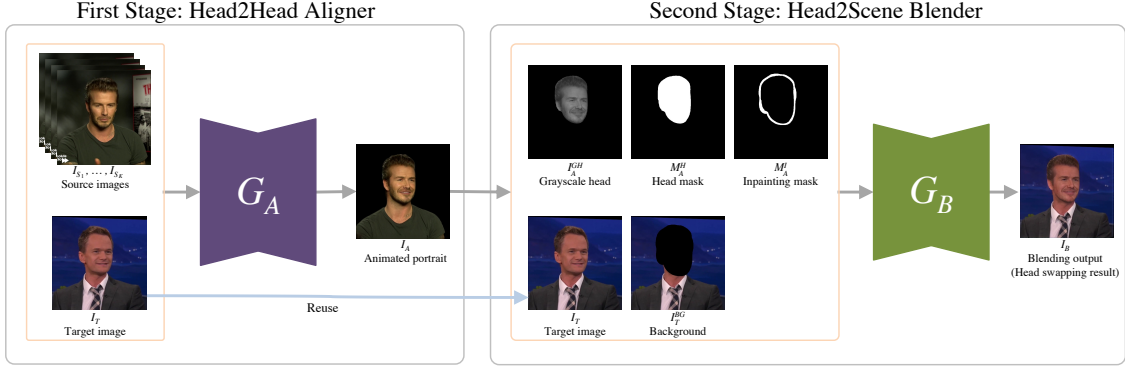


Figure 2. The overall architecture of our proposed pipeline for testing. First stage: the Head2Head Aligner G_A (detailed in Fig. 3) generates the animated portrait I_A , depicts the K -shot source image $I_{ID,k}$ ($k = 1, \dots, K$) at the same pose and expression of the target image I_T . Second stage: we concatenate grayscale head I_{GH} , head mask M_H , inpainting mask M_I , reference source I_T and background I_{BG} to go through the Head2Scene blender G_B (detailed in Fig. 5) to generate the blending output I_B (i.e., the head swapping result) in photo-realistic quality.

[10] solve the pose misalignment problem above. However, they tend to change the background color after the color transfer, making them unsuitable for head swapping where the background of the target must not be changed.

3. Method

In this work, we introduce the Head Swapper (HeSer), which is illustrated in Fig. 2. Let $I_S = \{I_{S_k} | k \in [1, K]\}$ denote the set of source images, and let I_T be the image of the target person. We aim to create a new blending output (i.e., head swapped result) I_B , with the head of S on the body and background of I_T . Notably, the pose, expression, and skin color of I_B should remain consistent with I_T , but the identity, head structure and hairstyle should be the same as S .

In this section, we first design a **Head2Head Aligner** (Sec. 3.1) to produce an *animated portrait* image I_A as an intermediate representation. Then I_B is derived by blending I_A with I_T in the **Head2Scene Blender** (Sec. 3.2).

3.1. Head2Head Aligner

Instead of directly building I_B , our first step is to produce an image which is positionally and emotionally aligned with I_T , but with the same identity as S . We name this intermediate *animated portrait* image I_A . We identify that all desired information required for I_A can be extracted from multi-scale perspectives on both the source and target, including: 1) the global head and hair structure information of I_S ; 2) the identity details on the faces of I_S ; 3) the global pose of I_T ; and 4) the detailed expressions of I_T .

To this end, we design the Head2Head Aligner as illustrated in Fig. 3. The Head2Head Aligner consists of a portrait encoder E_{por} , an identity encoder E_{id} , a pose encoder E_{pose} , an expression encoder E_{exp} , and a generator G .

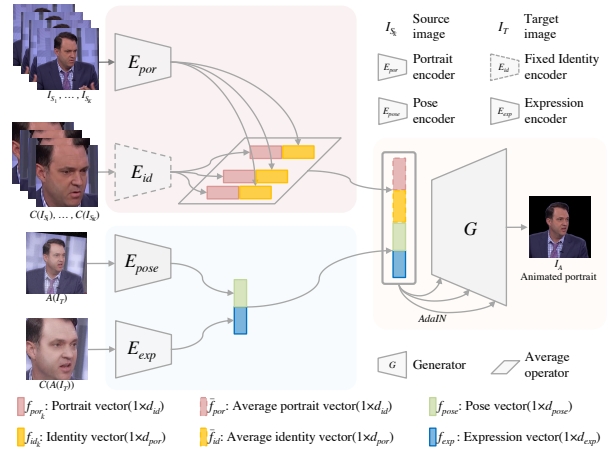


Figure 3. Overview of the proposed Head2Head Aligner for training. The average portrait vector \bar{f}_{por} , average identity vector \bar{f}_{id} , pose vector f_{pose} and expression vector f_{exp} are obtained from the portrait encoder, identity encoder, pose encoder and expression encoder respectively. Then the vectors above are sent into the animated portrait generation.

It is worth noting that during training, the source image and target image are obtained from the same video of a particular person (illustrated in Fig. 3), and they are derived from videos of different people during testing. Thus the information disentanglement should be taken into consideration. In the training stage, we randomly select $K + 1$ frames from the same video sequence of an individual person, where K frames are set as the K -shot source images I_S , while the remaining frame is set as the target image I_T . **Source Encoding.** Both the coarse and fine scales are leveraged for information encoding from the source and target. As for the source image, global portrait information including the head and hair, is directly extracted by an encoder




Figure 4. Illustration of data preprocessing for head2scene blender.

E_{por} to a d_{por} -dimensional vector $f_{por_k} = E_{por}(I_{S_k})$.

On the other hand, the identity encoder is adopted from a pretrained state-of-the-art face recognition model [5], which can provide more representative identity embedding [3, 17, 43]. More specifically, for each source image I_{S_k} , it first undergoes a central cropping transformation C , after which $C(I_{S_k})$ is passed through the identity encoder E_{id} to generate a d_{id} -dimensional identity embedding $f_{id_k} = E_{id}(C(I_{S_k}))$.

For randomly fetched K frames of source images $I_{S_1}, I_{S_2}, \dots, I_{S_K}$, a total of K portrait vectors f_{por_k} and K identity embeddings f_{id_k} are produced; we acquire \bar{f}_{por} and \bar{f}_{id} by taking the average of $\{f_{por_1}, \dots, f_{por_K}\}$ and $\{f_{id_1}, \dots, f_{id_K}\}$ respectively.

Target Encoding. To obtain both facial expression details and the global head pose, as shown at the bottom-left of Fig. 3, we apply the pose encoder E_{pose} and the expression encoder E_{exp} to extract the pose and detailed facial expression information respectively from the target image I_T . It first undergoes a series of augmentations [6] A before being sent into the pose encoder E_{pose} . An identity-agnostic d_{pose} -dimensional pose descriptor $f_{pose} = E_{pose}(A(I_T))$ is thus encoded. Then the augmented image is also processed with central cropping C mentioned above. The result $C(A(I_T))$ is fed into the expression encoder E_{exp} to produce a d_{exp} -dimensional expression vector f_{exp} .

Animated Portrait Generation. The feature descriptors containing information regarding the body, identity, pose and expression are then composed together to produce the *animated portrait*. Specifically, the features \bar{f}_{por} , \bar{f}_{id} , f_{pose} and f_{exp} are concatenated and processed by a generator via AdaIN [11] to produce the desired output I_A :

$$I_A = G(\bar{f}_{por}, \bar{f}_{id}, f_{pose}, f_{exp}) \quad (1)$$

The loss functions used are basically the reconstruction loss, perceptual loss, adversarial loss and identity loss, which are the same as [6]. Please refer to supplementary material for more details.

3.2. Head2Scene Blender

With the *animated portrait* image I_A above, our second step is to produce the blending output I_B . It should be a seamless combination of the head from I_A and the background and body in I_T . The challenges are 1) the head

shape on I_A could be significantly different from that of I_T , thus the blending procedure requires amending the “outliners”, *i.e.*, masking out the areas of the original head and fill in the gaps led by head shape mismatches; and 2) the color of the I_B should be consistent with I_T .

To this end, we design the Head2Scene Blender, which simultaneously solves both challenges as illustrated in Fig. 5. The key is to build “*color references*” which serve as guidance for both skin color transfer and background gap inpainting with a *Semantic-Guided Color Reference Creation* sub-module. Then the final results are generated through a *Blending UNet*. Below we will first introduce the necessary *Data preprocessing* steps for producing the *color references* and then introduces the sub-modules respectively.

Data preprocessing. The data preprocessing procedure prepares the materials for performing the color transfer and inpainting, which are binary masks that identifies the areas in need and the cropped images. Particularly, instead of directly modifying the animated head in I_A , we argue that it would be easier to reconfigure the color transfer problem into the problem of re-coloring its grey-scale version of the head I_A^{GH} only.

Detailedly, given the *animated portrait* I_A , we employ a parsing tool to obtain the segmentation masks M_A , which indicates the semantic areas of the portrait image. Then the head mask M_A^H of the target image is obtained by combining the masks belonging to the head area. Similarly, the segmentation masks M_T and head mask M_T^H of the target image I_T are obtained in the same way. The greyscale head can be computed as $I_A^{GH} = Gray(I_A * M_A^H)$.

Furthermore, the head mask M_T^H is dilated to an enlarged version \hat{M}_T^H , then the *target inpainting mask* $M_T^I = \hat{M}_T^H - M_T^H$ which will be further used for creating the *color reference* for inpainting. Different from the target inpainting mask above, the *animated inpainting mask* M_A^I is obtained via both M_A^H and M_T^H . More especially, the union set of M_A^H and M_T^H is dilated to an enlarged mask \hat{M}_A^H . Let the *animated inpainting mask* $M_A^I = \hat{M}_A^H - M_A^H$ denotes the area that needs inpainting when blending I_A with I_T . $I_T^{BG} = I_T * (1.0 - \hat{M}_A^H)$ is the background image without the target head.

Semantic-Guided Color Reference Creation. After the data processing steps, we need to re-color the greyscale head I_A^{GH} and fill the *animated inpainting mask* M_A^I . As stated before, our intuition is to provide a color reference for both of them, namely the *head-color reference* $I_{T \rightarrow A}^{HR}$ and the *inpainting reference* $I_{T \rightarrow A}^{IR}$, respectively.

Inspired by previous work [15, 37, 39], we identify that each *color reference* should spatially match the target region and provide pixel-wise color guidance. As the desired color should be consistent with the target image I_T , the color of the *references* could be derived by querying the coherent

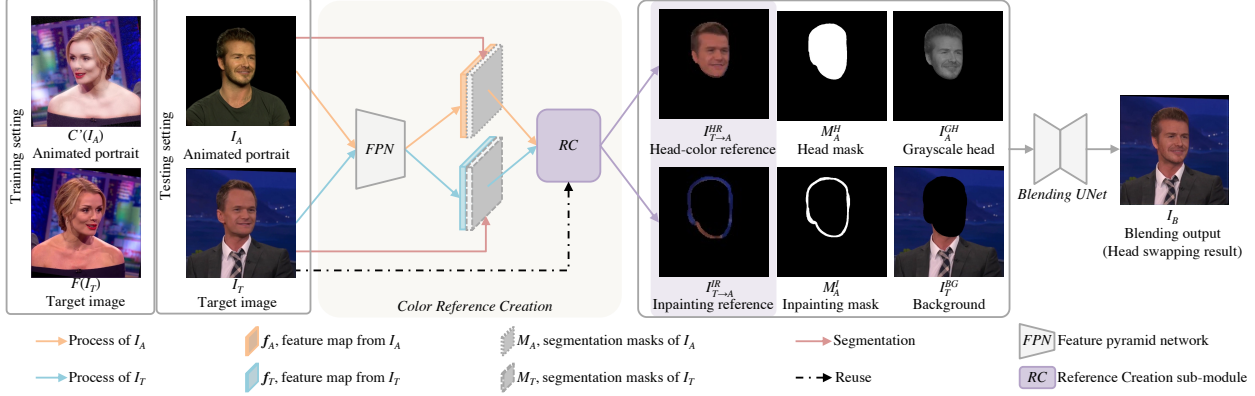


Figure 5. Illustration of our proposed Head2Scene Blender in the testing setting. The head-color reference $I_{T \rightarrow A}^{HR}$ and inpainting reference $I_{T \rightarrow A}^{IR}$ are first obtained using the Reference Creation sub-module. Then $I_{T \rightarrow A}^{HR}$, $I_{T \rightarrow A}^{IR}$, M_A^H , I_A^{BG} , M_A^I , and I_A^{GH} are fed into the blending UNet to produce the final result I_B . In the training phase, the same image is used for I_A , I_T and I_B .

color from I_T through correlation learning. This is achieved in the Reference Creation sub-module RC .

Due to the lack of paired data, the Reference Creation sub-module is still trained in a self-driven manner with I_A and I_T sourcing from the **same** image. As shown in the left-hand diagram of Fig. 5, in the training phase, the *animated portrait* I_A and target image I_T first undergo a random color augmentation C' and a random horizontal flip augmentation F respectively, which prevents the networks from directly leveraging the pixels from the same position. The augmented *animated portrait* $C'(I_A)$ and the random flipped *target image* $F(I_T)$ are then both sent into a feature pyramid network FPN [39] to produce the following semantic representation:

$$f_A = FPN(C'(I_A)). \quad (2)$$

$$f_T = FPN(F(I_T)). \quad (3)$$

The next step to colorization is to compute the correlations between extracted features on each spatial location, and resample pixel colors from the target image to the *color references*. Previous approaches [15, 37, 39] produce a prohibitively large memory footprint when estimating the correspondence due to the fact that the pairwise similarities are computed among all locations of the feature maps, despite their semantic independence. As the pixels in the image pairs with different semantic labels make little contribution to the correspondence matrix, we only compute the **correlations among the same semantic regions** individually. Our practice not only alleviates memory consumption but also avoids mismatched correlations.

Concretely, as shown in Fig. 6, for each **semantic region** $r \in \{face, hair, eye, nose, lip, tooth, inpainting\}$, we compute a correlation matrix $\Gamma^r \in \mathbb{R}^{N_A^r \times N_T^r}$, of which each element $\Gamma^r(u, v)$ is a pairwise feature correlation cal-

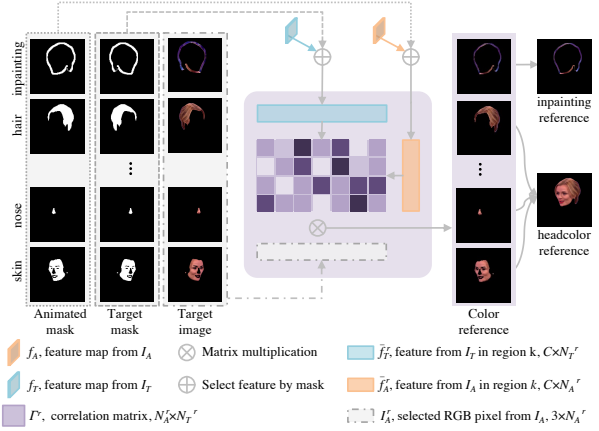


Figure 6. The illustration of our Semantic-Guided Color Reference Creation module, best viewed in color.

culated as:

$$\Gamma^r(u, v) = \frac{\tilde{f}_A^r(u)^T \tilde{f}_T^r(v)}{\|\tilde{f}_A^r(u)\| \|\tilde{f}_T^r(v)\|}, u \in M_A^r, v \in M_T^r, \quad (4)$$

where $\tilde{f}_A^r(u)$ and $\tilde{f}_T^r(v) \in \mathbb{R}^C$ represent the channel-wise centralized feature of f_A^r and f_T^r in positions u and v respectively, *i.e.*, $\tilde{f}_A^r(u) = f_A^r(u) - \text{mean}(f_A^r(u))$ and $\tilde{f}_T^r(v) = f_T^r(v) - \text{mean}(f_T^r(v))$, and the dimensions N_A^r and N_T^r are the pixel number of the semantic region r in the animated portrait image and target image respectively. Compared to the original correlation maps in prior works [15, 39], our approach reduces the computational complexity from $O(wh * wh)$ to $O(N_S^r * N_R^r)$, where w and h denote the spatial size of the original feature map.

As $\Gamma^r(u, v)$ indicates the similarity between $M_A(u)$ and $M_T(v)$, we normalize the $\Gamma^r(u, v)$ via softmax and compute their weighted contribution of these variables from the target image to the head and inpainting regions of the ani-

mated portrait image, as follows:

$$I_{T \rightarrow A}^r(u) = \sum_{v \in M_T^r} \text{softmax}_v(\Gamma^r(u, v)/\tau) \cdot I_T(v), u \in M_A^r. \quad (5)$$

where τ is the temperature coefficient. Finally, as shown in Fig. 6, the created inpainting reference $I_{T \rightarrow A}^{IR}$ is derived from I_T and M_T^I , while the warped head-color reference $I_{T \rightarrow A}^{HR}$ is obtained from I_T and the combination of M_T^r , where $r \in \{face, hair, eye, nose, lip, tooth\}$. We denote the overall estimation procedure as:

$$I_{T \rightarrow A}^{HR} = RC(f_A, M_A, f_T, M_T, I_T * M_T^H). \quad (6)$$

$$I_{T \rightarrow A}^{IR} = RC(f_A, M_A^I, f_T, M_T^I, I_T * M_T^I). \quad (7)$$

Blending UNet. After acquiring the created head-color/inpainting reference $I_{T \rightarrow A}^{HR}/I_{T \rightarrow A}^{IR}$, the head mask M_A^H , background I_T^{BG} , inpainting mask M_A^I and gray-scale head I_A^{GH} produced by the data processing, we concatenate them all channel-wisely, and pass them to a Blending UNet. Our goal is that: 1) the head-color distilled from the target image is transferred to the gray-scale head I_A^{GH} , where the identity is retained while the color is kept consistent with the remaining skin (such as neck) exposed in the rest of the reference body; 2) the missing region masked by the inpainting mask can be estimated. Consequently, the overall process of the Blending UNet B can be formulated as follows:

$$I_B = B(I_{T \rightarrow A}^{HR}, I_{T \rightarrow A}^{IR}, M_A^H, I_T^{BG}, M_A^I, I_A^{GH}). \quad (8)$$

Cycle Loss: The training relies on the reconstruction loss, perceptual loss, adversarial loss. Additionally, in order to guarantee that the warped headcolor/inpainting exemplars could learn a meaningful correspondence matrix, we introduce the cycle consistent loss.

$$L_c = \lambda_c \|I_{T \rightarrow A \rightarrow T} - I_T\|_1, \quad (9)$$

where $I_{T \rightarrow A \rightarrow T}$ is the color reference after cycled, and $I_{T \rightarrow A \rightarrow T}^k(u) = \sum_{v \in M_A^k} \text{softmax}_v(\Gamma^k(u, v)/\tau) \cdot I_{T \rightarrow A}(v), u \in M_T^k$. Besides, the additional target image I_T' coming from different image compared to I_A is also utilized to ensure the meaningful of warped exemplar:

$$L_{c'} = \lambda_c \|I_{T' \rightarrow A \rightarrow T'} - I_T\|_1. \quad (10)$$

Please refer to the supplementary materials for more details.

4. Experiments

Detailed descriptions of the data collection, along with more experimental details, are elaborated in the supplementary material.

Table 1. Average score from the user study, rating from 0 to 2.

Method	ID \uparrow	Exp \uparrow	Skin Color \uparrow	Inpainting \uparrow	Holistic \uparrow
FaceSwap [3]	0.488	1.208	1.264	–	0.492
Deepfacelab [23]	0.864	0.588	0.612	0.028	0.784
Ours	1.648	1.204	1.044	2.972	1.724

4.1. Headswap Results

We first evaluate our head swapping result. We compare our proposed HeSer with the state-of-the-art face swapping model [3] and Deepfacelab [23] as follows.

Qualitative comparison. The results in Fig. 7 demonstrate that our method can outperform other methods in multiple aspects, e.g., identity preservation, pose and expression consistency, skin color alignment, head-background coherence, and fidelity. The source image for face swapping is one of the 32-shot images, while the source images for Deepfacelab and our head swapping model are the same 32 frames. The face swapping method can swap the facial elements from source image to target image, and the mustache can be slightly transferred; however, the immutable face shape and hairstyle restrict the identity similarity in human cognition. Deepfacelab tends to animate poorly when there is no source image with a similar pose to target image; furthermore, it is incapable of inpainting the missing pixel and the performance of head-color alignment is significantly inferior to ours. For more vividly head swapping results generated by our method please refer to supplementary video.

Quantitative Comparison. We conduct a user study for quantitative comparison of our HeSer with the face swapping method [3] and Deepfacelab [23]. We ask the users to rank 1) how well the **ID** information is preserved in each method. 2) The emotion and pose similarity with the target image, we denote the results as **Exp**. 3) The consistency of the **Skin Color** between generated examples and the target torso. 4) The **Inpainting** smoothness between generated head and the background. 5) The **Holistic** quality of the generated frames.

We randomly select 25 groups of source images and 10 target images from the voxceleb2 test set, and a totally 250 results are obtained for each method. Then 15 different workers are asked to rate the randomly selected 25 results for every dimension on a scale of zero to two, where worst is zero and best is two. Table 1 shows the average score result from the user study. Our HeSer outperforms other methods by a large margin, except for the expression consistency and skin color alignment. Face swapping [3] directly injects the identity information to the target image thus naturally keeps better facial structures and skin colors.



Figure 7. Our head swapping results compared with other methods.

Table 2. Quantitative comparison on face reenactment.

Method	$E_{ID} \downarrow$	$E_P \downarrow$	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow
FOMM [27]	0.49	0.275	0.76	0.18	30.92
LPD [6]	0.71	0.063	0.52	0.50	28.84
Siarohin et al [28]	0.71	0.137	0.73	0.20	30.01
Ours	0.53	0.026	0.77	0.19	31.33
LPD (32-shot + ft)	0.23	0.024	0.62	0.36	29.46
Ours (32-shot + ft)	0.22	0.024	0.89	0.12	33.26

4.2. Superiority of the Head2Head Aligner

In this subsection, we evaluate Head2Head Aligner independently. As the Head2Head Aligner is devised to animate the source image with the pose and expression of the target image, which is similar to the task of reenactment, different reenactment strategies are compared. The performance is discussed from two perspectives: K-shot inputs, and fine-tuning or not. As some of our 2D-based animated competitors (FOMM [27] and Siarohin et al. [28]) do not support few-shot inputs and fine-tuning, we split our evaluations into 1) a one-shot setting, where all competitors are compared; 2) otherwise, only LPD is compared.

Evaluation Metrics. Five metrics are included in terms of head reenactment evaluation: (1) **Identity error** (E_{ID})




Figure 8. Qualitative cross-id animated results of 32-shot fine-tuned model with 600 iterations. Source, target, and our animated result are re-aligned following the instruction of LPD.

using the cosine similarity between face embeddings. (2) **Pose reconstruction error** (E_P), measuring the pose error between the synthesized and ground truth images as performed in [6, 35]. (3) Pixel-wise Reconstruction fidelity using **PSNR**. (4) Semantic perceptual similarity via the AlexNet-based **LPIPS** metric [40]. (5) The perceived quality via the structural similarity index measure (**SSIM**) [32].

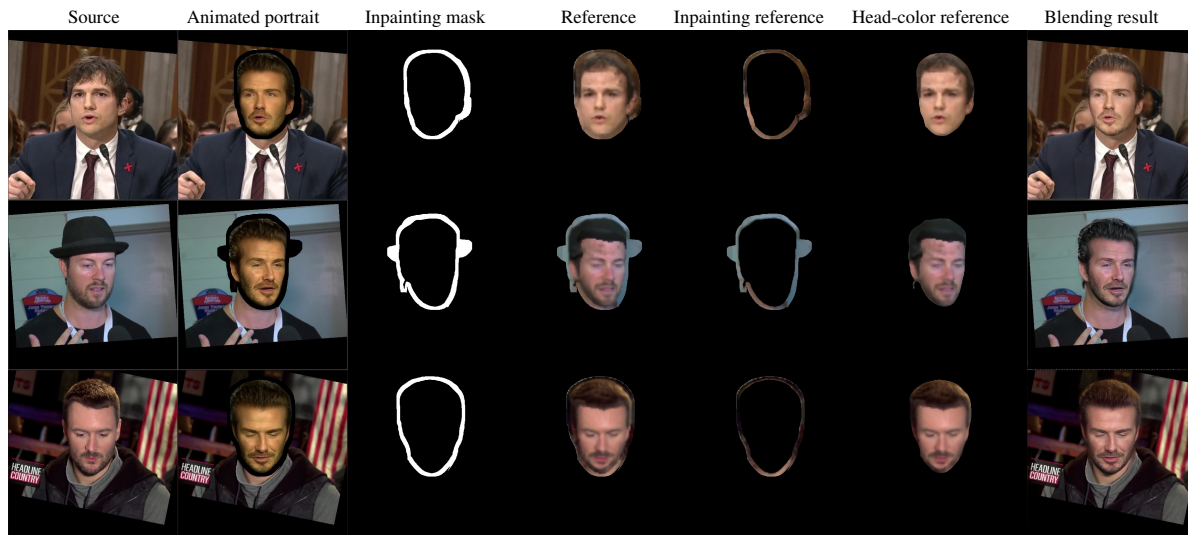


Figure 9. Illustration of head-color/inpainting reference in our Head2Scene Blender. Best viewed in color.

One-shot setting Evaluation. We first compare the performance of our animated method with four competitors via five evaluation metrics under the one-shot face reenactment setting. The quantitative results are shown in Table 2. Our method substantially outperforms all other animated models on every metric, except for the identity error and LPIPS, which is slightly lower than FOMM. For more qualitative comparison results please refer to supplementary material.

Impact of K-shot and fine-tuning. We further analyze the effect of increasing the K-shot number with subject-specific fine-tuning or a meta-learned model, which is the setting illustrated in Fig. 1. Specifically, we set the model that takes 32 frames as input and leverage 600 iterations of finetuning as the standard setting. The quantitative comparisons are illustrated in Table 2. More quantitative evaluations on the influences of shot number and finetuning iterations are plotted in supplementary material. It shows that our method almost outperforms other methods on all metrics with various settings. The performances of all metrics tend to improve as the k-shot number increases and fine-tuning is added respectively, while the identity error and pose reconstruction error of our animated method exhibit prominent improvement relative to LPD, demonstrating that our proposed scheme is capable of generating animated portraits with smaller identity gap and higher pose consistency.

It is further noteworthy that our animated scheme outperforms LPD with only about one-fifth amounts of the training data in LPD. Though the identity error of our method and LPD is almost equal when the number of tuning steps is sufficiently large (such as 600), the facial expressions and emotions of our animated portraits are significantly better than those of LPD, as illustrated in Fig. 8. Please refer to supplementary material for more detailed qualitative results.

4.3. Effectiveness of Head2Scene Blender

The representative blending results are shown in the last column of Fig. 9. Our Head2Scene Blender can produce head swapping portraits of photo-realistic quality. The head-color reference (sixth column in Fig. 9) is created from the target image, while the head structure is consistent with the animated head. Similar behavior is observed in the inpainting reference (fifth column in Fig. 9), where the estimated region style is faithful to the reference background. It is noteworthy that the hat region of the inpainting reference in the second row is constructed from the source background, which further demonstrates the inpainting effectiveness of our blending module.

5. Conclusion and Discussion

Conclusion. In this paper, we propose the Head Swapper (Heser), which achieves few-shot head swapping for in-the-wild scenarios. Specifically, our Head2Head Aligner generates high-fidelity reenact results with high pose and expression consistency, and the Head2Scene Blender seamlessly blends the aligned source head to the target image while maintaining the color of the target person. Extensive experiments demonstrate that the HeSer can achieve superior head swapping results on a variety of scenes.

Broader Impact. Vivid video synthesis technologies create possibilities for immoral behaviors. Recent generative models have greatly influenced identity safe, image authenticity, etc. We will share the results of HeSer to the face/head forgery detection community for the healthy development of the AI technology.

Acknowledgement. This work is supported by the CCF-Baidu Open Fund.

References

- [1] Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, and Gang Hua. Towards open-set identity preserving face synthesis. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 2
- [2] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. Pairedcyclegan: Asymmetric style transfer for applying and removing makeup. *cvpr*, 2018. 2
- [3] Renwang Chen, Xuanhong Chen, and Yanhao. Ge. Simswap: An efficient framework for high fidelity face swapping. In *ACM Int. Conf. Multimedia*, 2020. 1, 2, 4, 6, 7
- [4] Zhuo Chen, Chaoyue Wang, Bo Yuan, and Dacheng Tao. Puppeteergan: Arbitrary portrait animation with semantic-aware appearance transformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [5] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 4
- [6] Burkov Egor, Pasechnik Igor, Grigorev Artur, and Lempitsky Victor. Neural head reenactment with latent pose descriptors. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2, 4, 7
- [7] Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. *arXiv preprint arXiv:2012.03065*, 2020. 2
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [9] Qiao Gu, GuanZhi Wang, Tik Chiu Mang, YuWing Tai, and ChiKeung Tang. Ladrn: Local adversarial disentangling network for facial makeup and de-makeup. *iccv*, 2019. 2
- [10] Han, Chu Deng, Hongmin Han, Guoqiang cai, Shengfeng Han, and He. Spatially-invariant style-codes controlled makeup transfer. In *CVPR*, 2021. 3
- [11] Xun Huang and Serge. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE Conf. Conference of the International Speech Communication Association.*, 2017. 4
- [12] Huiwen, Jingwan Chang, Fisher Lu, Adam Yu, and Finkelstein. Pairedcyclegan: Asymmetric style transfer for applying and removing makeup. In *CVPR*, 2018. 2
- [13] Hungjen, Kaming Chen, Szuyu Hui, Liwu Wang, Honghan Tsao, Wenhuan Shuai, and Cheng. Beautyglow: On-demand makeup transfer framework with reversible generative network. In *CVPR*, 2019. 2
- [14] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Wayne Wu, Chan Change Loy, Xun Cao, and Feng Xu. Audio-driven emotional video portraits. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [15] Wentao Jiang, Si Liu, Chen Gao, Jie Cao, Ran He, Jiashi Feng, and Shuicheng Yan. X2face: A network for controlling face generation by using images, audio, and pose codes. *cvpr*, 2020. 2, 4, 5
- [16] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *SIGGRAPH*, 2018. 2
- [17] Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, and Fang. Wen. Faceshifter: Towards high fidelity and occlusion aware face swapping. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 1, 4
- [18] Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, and Fang. Wen. Faceshifter: Towards high fidelity and occlusion aware face swapping. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 2
- [19] Borong Liang, Yan Pan, Zhizhi Guo, Hang Zhou, Zhibin Hong, Xiaoguang Han, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. Expressive talking head generation with granular audio-visual control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [20] Yuanxun Lu, Jinxiang Chai, and Xun Cao. Live speech portraits: real-time photorealistic talking-head animation. *ACM Transactions on Graphics (TOG)*, 40(6):1–17, 2021. 2
- [21] Yuval Nirkin, Yosi Keller, and Tal Hassner. Fsgan: Subject agnostic face swapping and reenactment. *iccv*, 2019. 1, 2
- [22] Yuval Nirkin, Iacopo Masi, Anh Tran Tuan, Tal Hassner, and Gerard Medioni. On face segmentation, face swapping, and face perception. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 98–105. IEEE, 2018. 2
- [23] Ivan Perov, Daiheng Gao, Nikolay Chervoniy, Kunlin Liu, Sugasa Marangonda, Chris Um, Mr Dpfks, Carl Shift Facenheim, Luis RP, Jian Jiang, et al. Deepfacelab: Integrated, flexible and extensible face-swapping framework. *arXiv preprint arXiv:2005.05535*, 2020. 1, 2, 6, 7
- [24] Kripasindhu Sarkar, Vladislav Golyanik, Lingjie Liu, and Christian Theobalt. Style and pose control for image synthesis of humans from a single monocular view. *arXiv preprint arXiv:2102.11263*, 2021. 1, 2
- [25] Si, Wentao Liu, Chen Jiang, Ran Gao, Jiashi He, Bo Feng, Shuicheng Li, and Yan. Psgan++: Robust detail-preserving makeup transfer and removal. In *PAMI*, 2021. 2
- [26] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. *CVPR*, 2019. 2
- [27] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *neurips*, 2019. 2, 7
- [28] Aliaksandr Siarohin, Oliver Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. Motion representations for articulated animation. In *CVPR*, 2021. 2, 7
- [29] Tero, Samuli Karras, Miika Laine, Janne Aittala, Jaakko Hellsten, Timo Lehtinen, and Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 2
- [30] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Trans. Graph.*, 2019. 2
- [31] Tingting, Ruihe Li, Chao Qian, Si Dong, Qiong Liu, Wenwu Yan, Liang Zhu, and Lin. Beautygan: Instance-level facial makeup transfer with deep generative adversarial network. In *ACMMM*, 2018. 2

- [32] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. In *IEEE transactions on image processing*, 2004. 7
- [33] Zhiliang Xu, Zhibin Hong, Changxing Ding, Zhen Zhu, Junyu Han, Jingtuo Liu, and Errui Ding. Mobilefaceswap: A lightweight framework for video face swapping. *AAAI*, 2022. 2
- [34] Zhiliang Xu, Xiyu Yu, Zhibin Hong, Zhen Zhu, Junyu Han, Jingtuo Liu, Errui Ding, and Xiang Bai. Facecontroller: Controllable attribute editing for face in the wild. *AAAI*, 2021. 2
- [35] Egor Zakharov, Aleksei Ivakhnenko, Aliaksandra Shysheya, and Victor Lempitsky. Fast bi-layer neural synthesis of one-shot realistic head avatars. In *Eur. Conf. Comput. Vis.*, August 2020. 7
- [36] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *Int. Conf. Comput. Vis.*, October 2019. 2
- [37] Bo Zhang, Mingming He, Jing Liao, V. Sander Pedro, Lu Yuan, Amine Bermak, and Dong Chen. Deep exemplar-based video colorization. In *CVPR*, pages 8052–8061, 2019. 4, 5
- [38] Jiangning Zhang, Xianfang Zeng, Mengmeng Wang, Yusu Pan, Liang Liu, Yong Liu, Yu Ding, and Changjie Fan. Freenet: Multi-identity face reenactment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [39] Pan Zhang, Bo Zhang, Dong Chen, Lu Yuan, and Fang Wen. Cross-domain correspondence learning for exemplar-based image translation. In *CVPR*, pages 5143–5153, 2020. 4, 5
- [40] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 586–595, 2018. 7
- [41] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [42] Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. Talking face generation by adversarially disentangled audio-visual representation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2019. 2
- [43] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 4
- [44] 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 (CVPR)*, pages 4834–4844, 2021. 2