

FaceController: Controllable Attribute Editing for Face in the Wild

Zhiliang Xu ^{1*}, Xiyu Yu ^{2*}, Zhibin Hong ², Zhen Zhu ¹, Junyu Han ²,
Jingtuo Liu ², Errui Ding ², Xiang Bai ^{1†}

¹ Huazhong University of Science and Technology

² Baidu Inc.

{zhiliangxu1, zzhu, xbai}@hust.edu.cn, {yuxiyu, hongzhibin, hanjunyu, liujingtuo, dingerrui}@baidu.com

Abstract

Face attribute editing aims to generate faces with one or multiple desired face attributes manipulated while other details are preserved. Unlike prior works such as GAN inversion, which has an expensive reverse mapping process, we propose a simple feed-forward network to generate high-fidelity manipulated faces. By simply employing some existing and easy-obtainable prior information, our method can control, transfer, and edit diverse attributes of faces in the wild. The proposed method can consequently be applied to various applications such as face swapping, face relighting, and makeup transfer. In our method, we decouple identity, expression, pose, and illumination using 3D priors; separate texture and colors by using region-wise style codes. All the information is embedded into adversarial learning by our identity-style normalization module. Disentanglement losses are proposed to enhance the generator to extract information independently from each attribute. Comprehensive quantitative and qualitative evaluations have been conducted. In a single framework, our method achieves the best or competitive scores on a variety of face applications.

Introduction

Manipulating and editing real faces in the wild has countless applications in the areas of visual effects and e-commerce. Face editing tasks often benefit from the highly disentangled latent representation of face attributes such that one can precisely manipulate corresponding attributes. For example, face swapping, a special case of face editing, aims to preserve the identity of a source face while retaining the pose and expression of a target face. Therefore, achieving fortified disentanglement of face attribute latent codes is the core of many works (Zhang et al. 2019; Shen et al. 2020; Ha et al. 2019; Zeng et al. 2020).

However, deriving such disentangled latent representation from a real face is still challenging. Directly mapping a face into the latent space by an encoder (Chen et al. 2018; Klys, Snell, and Zemel 2018) can be efficient. But it is not ensured to extract latent codes that express semantically-meaningful factors of variations (Nie et al. 2020). Thus, another common practice is to obtain the reverse mapping in latent space

for a real face, known as *GAN inversion*. For example, (Abdal, Qin, and Wonka 2019; Zhu et al. 2020a) opt to embed a real face into the latent space of StyleGAN (Karras, Laine, and Aila 2019) by optimizing a randomly initialized code through back-propagation. In practice, approaches of this kind are often time-consuming and cumbersome.

Inspired by recent face synthesis methods such as (Gecer et al. 2018; Deng et al. 2020), we propose a simple feed-forward face generation network that is fed with some existing and easy-obtainable prior information. As such, this method can avoid the costly learning process of disentangled representation. In this paper, we name our method as FaceController, which enables the control of face attributes. We embed 3DMM coefficients into an adversarial learning framework. However, we observe there exists a domain gap between real faces and the counterpart ones rendered from 3DMM. For example, the makeup of a woman cannot be perfectly expressed by 3DMM. To combat this, we further employ the embedding of identity and region-wise style codes as side information to complement the information gap between 3DMM coefficients and real faces, making our method stand out from previous works.

In the proposed FaceController, a new identity-style normalization module is proposed to fuse all the prior information. The identity-style normalization module is an extension of SPADE (Park et al. 2019) in which we use the embedding of identity and region-wise style codes to modulate the feature maps originated from 3DMM coefficients. Furthermore, to make the generator generalize to different editing, we devise unsupervised face generation losses on identity, landmarks, and textures to enforce the generator to control each attribute independently.

FaceController provides more comprehensive functions of real face editing by precisely controlling the latent codes of a wide range of face attributes. Unlike previous methods, which mainly address the disentanglement of attributes such as identity and expression, identity and pose, our method is able to disentangle in total five different attributes, including identity, pose, expression, illumination, and local region styles. As a result, FaceController enables a variety of face applications. For example, it can transfer the lip color or eye shadow makeup by editing the region-wise codes; it can also be applied to adjust the illumination of a given face image. In this paper, we empirically validate that the proposed

* Authors contribute equally.

† Corresponding author.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

method achieves better quantitative scores and visual results in most of the above applications.

The main contributions of this paper are: (1) FaceController generates high-quality face with desired attributes using a simple feed-forward face generation network and improves efficiency by employing easily-obtainable prior information rather than the time-consuming reverse mapping process; (2) FaceController also increases the freedom of disentangled and controllable dimensions of attributes, making it better to control, transfer, and edit face attributes.

Related Work

Disentangled representation learning. In face attribute editing, many attempts have been made to learn an independent representation of face attributes. Early attempts include unsupervised disentangled latent representation learning of real face images. For example, InfoGAN (Chen et al. 2016) and its variants (Lin et al. 2019) target to maximize the mutual information between latent representation. HoloGAN (Nguyen-Phuoc et al. 2019) learns 3D representation from natural images to disentangle identity and pose. StyleGAN (Karras, Laine, and Aila 2019) can synthesize highly photorealistic face images and allow scale-specific modifications to the styles. Variational Autoencoder (VAE) based methods (Shen and Liu 2017; Razavi, van den Oord, and Vinyals 2019; Chen et al. 2018; Sun et al. 2018) often encourage independence of the latent representation by regularizing their total correlation. These methods achieved a promising performance. However, the disentangled latent codes are not provably identifiable (Locatello et al. 2019), and strong regularization on latent representation sometimes may be harmful to generate high-fidelity face images due to information loss (He et al. 2019).

Other attempts aim to extract disentangled latent representation using some explicit or implicit supervision information (Bouchacourt, Tomioka, and Nowozin 2018). For example, (Deng et al. 2020; Tewari et al. 2020) aim to manipulate the StyleGAN latent space by using 3DMM information as supervision during learning. (Shen et al. 2020) tries to classify the latent space into multiple linear subspaces with semantic meaning. It is also possible to impose explicit supervision on latent codes (Nie et al. 2020). Most of these works can manipulate fake faces with high precision. However, in order to edit real faces, an optimization-based embedding process named GAN inversion has to be conducted to infer the best latent code for this face. It is often time-consuming.

Facial manipulation using 3D priors. 3DMMs (Banz, Vetter et al. 1999; Cao et al. 2013) represent the shape and appearance of faces by projecting them into low-dimensional spaces using techniques such as Principal Component Analysis (PCA). 3DMMs provide an explicit way to independently represent information such as identity, expression, and pose. Therefore, it can be naturally applied to face manipulation. For example, Face2Face (Thies et al. 2016) models both the driving and source face via a 3DMM; then applies the expression components of driving face to source face. (Kim et al. 2018; Xu et al. 2020) try to manipu-

late expression and pose by modifying 3DMM coefficients. (Gecer et al. 2018) proposes a semi-supervised adversarial framework for generating photorealistic faces directly from 3DMM. However, we observe that simply deploying 3D priors often results in a domain gap between real faces and generated faces due to the information loss of 3DMM. To minimize this domain gap, we propose an identity-style normalization module to fuse some side information. Consequently, we can obtain more photorealistic faces.

Controllable Attribute Editing

In this paper, we aim to efficiently edit faces in the wild with desired attributes. We propose a simple feed-forward framework to generate the edited faces directly from some easily obtainable latent representation of attributes. Specifically, denote I as a face image and $\{a_1, \dots, a_n\}$ as its latent representation of n different attributes. With simply replacing a_i , ($i \in 1, \dots, n$) by desired latent code a'_i , we can generate desired face image I' and achieve attribute editing.

To achieve this, FaceController requires at least three stages: 1) an encoding stage that aims to produce attribute-independent latent codes; 2) an attribute editing stage that modifies the produced latent codes to the desired ones through simple replacement; 3) a decoding stage that receives the edited latent codes as input and produces face images that correspond to the edited latent codes.

FaceController has specific and useful designs for stronger disentanglement of a wide range of face attributes and interpreting them to high-quality and plausible face images that conform to the edited latent codes. Taking face swapping as an example, the architecture of FaceController generator is shown in Fig. 1 (a). Each part of FaceController is built with clear and reasonable motivations and narrated as follows.

Face Attribute Disentanglement

Among available solutions to perform face attribute disentanglement, 3DMM (Egger et al. 2020) stands as a near-optimal option because it disentangles a wide range of different attributes including identity, expression, pose, illumination, and texture. Many previous works (Deng et al. 2020) conduct attributes disentanglement based on 3DMM coefficients and show impressive results. However, we notice that the expressive power of its *identity* and *texture* coefficients are not sufficient to reconstruct a perceptually photo-realistic face. This drawback attributes to the limited volume of 3D face data where the original 3DMM (Deng et al. 2019b) is trained, limiting its generalization ability to real-world faces. Considering this, we believe identity embedding and region-wise texture representation should be included to complement the insufficient 3DMM identity and texture representation. A recent work (Zhu et al. 2020b) particularly centers on region-wise face control while it is not designed for other attributes control. To successfully modulate 3DMM attributes control and region-wise styles control in a unified architecture, we first extract 3DMM coefficients and then build two encoders that encode more representative identity and region-wise styles, respectively.

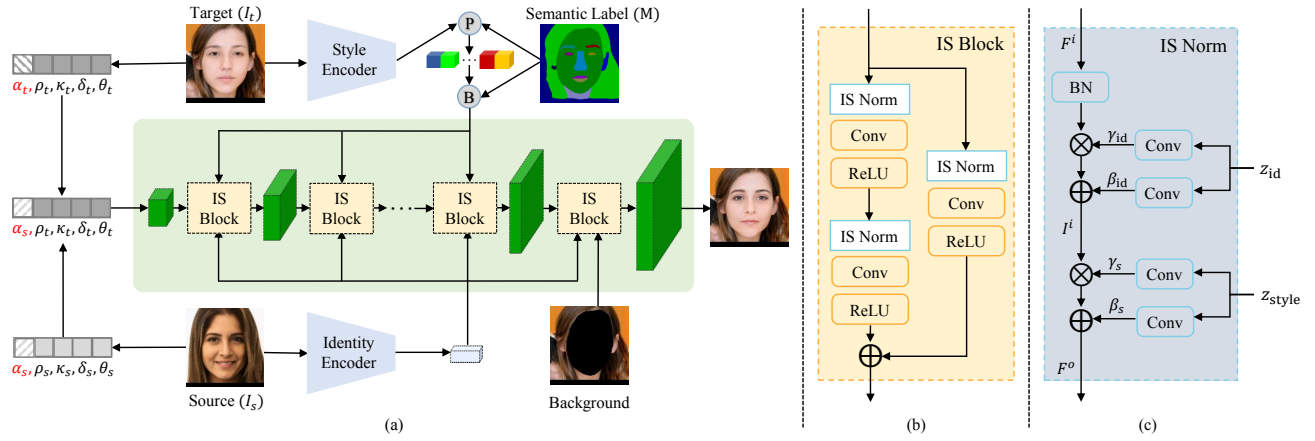


Figure 1: (a) The architecture of FaceController. We take the task of face swapping as an example. The 3DMM coefficients $(\alpha_s, \rho_s, \kappa_s, \delta_s, \theta_s)$ and $(\alpha_t, \rho_t, \kappa_t, \delta_t, \theta_t)$ are extracted from source and target images I_s and I_t , respectively. First, we transfer the coefficients of identity from source to target and obtain $(\alpha_s, \rho_t, \kappa_t, \delta_t, \theta_t)$. The transferred 3DMM coefficients with the style and identity codes are fed into the generator to synthesis a swapped face. Here, P denotes region-wise pooling; and B denotes the broadcasting process which is the inverted operation of pooling. (b) The architecture of IS block. IS block contains several Identity-Style Normalization modules that integrate all the information of above latent representation. (c) The architecture of IS normalization module. z_{id} and z_{style} represent identity and region-wise style respectively.

3DMM coefficients extraction. Specifically in a 3DMM, a face is often modeled with 3D shape \mathcal{S} and texture \mathcal{T} , which is obtained by:

$$\begin{aligned} \mathcal{S} &= \bar{\mathcal{S}} + \alpha I_{base} + \rho E_{base}, \\ \mathcal{T} &= \bar{\mathcal{T}} + \delta T_{base}. \end{aligned} \quad (1)$$

In the above equation, $\bar{\mathcal{S}}$ and $\bar{\mathcal{T}}$ are the average face shape and texture, respectively; I_{base} , E_{base} , and T_{base} are the PCA bases of identity, expression, and texture, respectively; $\alpha \in \mathbb{R}^{80}$, $\rho \in \mathbb{R}^{64}$, and $\delta \in \mathbb{R}^{80}$ are the corresponding 3DMM coefficients for rendering a 3D face. Besides shape and texture, the illumination $\kappa \in \mathbb{R}^{27}$ of the face can be approximated with Spherical Harmonics (SH) and the pose is defined as three rotation angles. For simplicity, we use θ to denote the parameters of rotation and translation. Luckily, $(\alpha, \rho, \kappa, \delta, \theta)$ can be obtained by an off-the-shelf 3D face reconstruction network (Deng et al. 2019b) and we use this network as the 3DMM coefficients extractor.

Identity encoder. As discussed previously, 3DMM is not able to perfectly depict the identity of a wild face. To complement the identity information, we employ an identity encoder to extract identity information from real images. Here, a state-of-the-art pretrained face recognition model (Deng et al. 2019a) is used as the identity encoder. We use the last feature maps before the final fully-connected (FC) layer to obtain a precise and high-level representation of identity. It is worthy to note that the face images fed into the 3DMM coefficients extractor and identity encoder are aligned in a different way. Then, in our identity encoder, we introduce the spatial transformation network (Jaderberg et al. 2015) to provide a precisely aligned face for our identity encoder.

Region-wise style encoder. To support local face region editing, we aim to obtain region-wise style codes by taking

advantage of the semantic segmentation of face images. A segmented region that is assigned with the same class label associates with a style code. We adopt the encoder of SEAN (Zhu et al. 2020b) as our region-wise style encoder for its superior ability to extract region-wise styles. Specifically, a region pooling layer is applied at the end of the encoder for each semantic label to extract a 512 dimensional style code. This region-wise style encoder naturally provides a disentangled representation of style for local face regions.

Face Attribute Remapping

At this point, we have acquired different sorts of latent codes—3DMM coefficients, identity embedding, and region-wise style codes—for detailed and precise attribute disentanglement. The next is how to leverage these codes and decode them back to a plausible generated face image.

Our task is generally related to translate semantic labels to natural face images with specific region-wise styles. SPADE (Park et al. 2019) is especially useful to this task and some recent works SEAN (Zhu et al. 2020b) and SMIS (Zhu et al. 2020c) have built their methods based on SPADE to support region-wise editing. Based on this observation, we also take the spirit of SPADE to design our decoder. However, our task requires the decoder to additionally consider region-irrelevant attributes such as identity. To support this, we design an *Identity-Style Normalization module* that integrates the identity and region-wise style information to the decoder, shown in Fig. 1 (c).

Let $F^i \in \mathbb{R}^{B \times C^i \times H^i \times W^i}$ denote the feature maps of the i -th layer of the decoder; B , C^i , H^i , and W^i represent the batch size, channel numbers, height, and width, respectively. We first perform the batch normalization on F^i and get \bar{F}^i . To integrate identity information, we map the identity embedding to a vector z_{id} with the same dimension of channel

numbers by using an FC layer. After that, two normalization parameters, γ_{id} and β_{id} , are learned from this vector to modulate \overline{F}^i like a style code. The modulated output is,

$$I^i = \gamma_{id} \circ \overline{F}^i + \beta_{id}, \quad (2)$$

we denote \circ as the element-wise production.

To further integrate region-wise style codes, we first extract a style matrix $S \in \mathbb{R}^{B \times N \times D}$ for region style codes by our region-wise style encoder. We also provide a corresponding segmentation map $M^i \in \mathbb{R}^{B \times N \times H^i \times W^i}$ which is interpolated from the segmentation map of the input face. Here, N denotes the number of semantic classes; D is the dimension of style code for each semantic class. Such segmentation map has N channels as every channel is responsible for a single class. To adaptively adjust a specific local region by its corresponding style code, we broadcast the style code into a feature map according to its segmentation mask. This broadcast process can be achieved by applying a batch-version matrix multiplication between the segmentation map M^i and style matrix S , i.e.,

$$z_{style} = S^T \times M^i, \quad (3)$$

where $S^T \in \mathbb{R}^{B \times D \times N}$ denotes the transposed matrix of S ; $z_{style} \in \mathbb{R}^{B \times D \times H^i \times W^i}$ are the broadcast style codes. Similar to identity normalization, we learn two parameters γ_s and β_s from z_{style} to modulate I^i , then we have the output of IS normalization module,

$$F^o = \gamma_s \circ I^i + \beta_s. \quad (4)$$

Training Losses

During training, we mainly consider two different training processes, i.e., face reconstruction and unsupervised face generation training. During face reconstruction training, the model receives attribute latent codes from the same face image and tries to reconstruct such a face. In this case, I_s is equal to I_t as depicted in Fig. 1. Such training ensures that the generated image is reasonable when $I_s = I_t$. However, the model is not ensured to work well when $I_s \neq I_t$. Consequently, we include unsupervised face generation training to guarantee even with unpair input ($I_s \neq I_t$), the model is still able to translate the combined latent codes to a plausible face image. Note that GAN loss is applied to both training processes for better image fidelity.

Face Reconstruction Loss. Following FaceShifter (Li et al. 2020), we utilize 20% training data for face reconstruction. We implement face reconstruction loss using the perceptual loss (Johnson, Alahi, and Fei-Fei 2016) to supervise the reconstruction. The intermediate feature maps of the generated face I_g and the given face I_t are extracted from a pre-trained VGG network (Simonyan and Zisserman 2014), and then we conduct pixel-level reconstruction as follows,

$$\mathcal{L}_{per} = \frac{1}{2} \|\mathcal{F}_{per}(I_t) - \mathcal{F}_{per}(I_g)\|^2, \quad (5)$$

where \mathcal{F}_{per} denotes the extractor of feature maps.




Figure 2: Demonstration of our landmark loss. Our aligned landmarks can better preserve the identity of the source image. (a) Source face with its landmarks. (b) Target face with its landmarks. (c) The swapped face with the landmarks of target face. As we can see, the landmarks and the face image are not well aligned (Zoom view better). (d) The swapped face with our aligned landmarks.

Unsupervised Face Generation Loss. The main purpose of our work is to support free and dynamic face attribute control over face images. Basically, when we alter some attributes, we wish to see results with these attributes changed accordingly while other attributes are unchanged. We design *identity loss*, *landmark loss* and *histogram matching loss* to assist the unsupervised training and further fortify attribute disentanglement.

If we aim to transfer the attributes such as expression and pose from I_t to I_s while keeping the identity, we can encourage the generated face to keep the same identity with I_s by utilizing the identity encoder:

$$\mathcal{L}_{id} = 1 - \cos(\mathcal{F}_{id}(I_g), \mathcal{F}_{id}(I_s)). \quad (6)$$

Here, \mathcal{L}_{id} refers to the identity loss and \mathcal{F}_{id} is the extractor of identity embedding. We use cosine similarity to estimate the similarity between the identity embedding of the generated image and the source image.

Then, if we want to change attributes like identity and textures from I_s to I_t but retain the expression and pose, we can leverage the landmark loss to ensure expression and pose consistency between I_t and the generated face I_g :

$$\mathcal{L}_{lm} = \frac{1}{2} \|\mathcal{F}_{lm}(I_g) - \mathcal{F}_{lm}(I_t)\|^2, \quad (7)$$

where \mathcal{F}_{lm} represents the face landmark extractor.

In this paper, the landmark loss is also specifically designed for identity preservation. As we want to keep the generated image I_g having the same expression and pose with target image I_t , a common implementation is to constrain the I_g and I_t to have the same landmarks. However, face landmarks contain the shape of the eyes, mouth, eyebrows, and nose. These are also linked with identity information. As shown in Fig. 2 (a) and Fig. 2 (b), different persons have different face features and landmarks. Consequently, as demonstrated in Fig. 2 (c), using landmarks of I_t to constrain the generated image is not a good way as we want to preserve identity information of I_s . We need to adjust the landmarks to have the same pose and expression with I_t while keeping the same identity information with I_s . To solve this problem, we utilize the swapped 3DMM coefficients as shown in Fig. 1 to extract 3D landmarks. These landmarks often can precisely keep the same face features with I_s , seen in Fig. 2

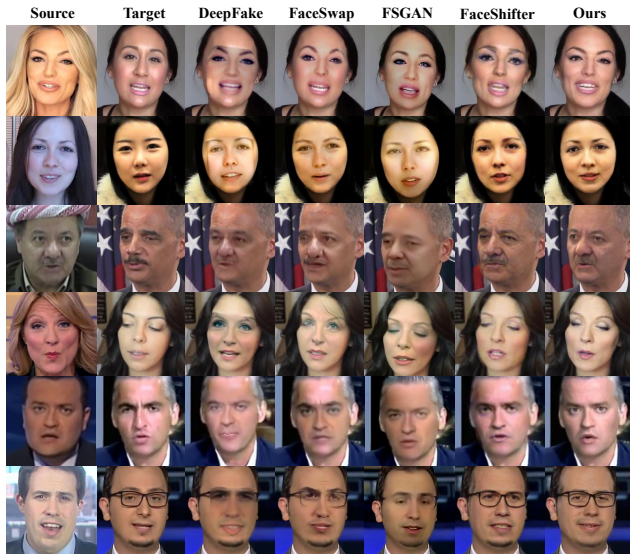


Figure 3: Comparison with DeepFake, FaceSwap, FSGAN, and FaceShifter on FaceForensics++ dataset.

(d). Then, matching the 3D landmarks of I_g and the aligned landmarks can better preserve the identity information.

In order to keep the style of target image, we further introduce a histogram matching loss (Risser, Wilmot, and Barnes 2017) to encourage region-wise style consistency with a target face. Let $I_{re} = \text{HM}(I_t, I_g)$ be a remapped image obtained by the histogram matching between I_g and I_t . Then, we have

$$\mathcal{L}_{\text{hm}} = \frac{1}{2} \|I_g - I_{re}\|^2. \quad (8)$$

For this loss to be minimized, the color distribution of I_g is learned to be similar to I_t , making local styles of the generated face more similar to the target.

Full Loss During training, the full loss is defined as a weighted sum of above losses:

$$\mathcal{L} = \mathcal{L}_{\text{adv}} + \lambda_{\text{id}} \mathcal{L}_{\text{id}} + \lambda_{\text{lm}} \mathcal{L}_{\text{lm}} + \lambda_{\text{hm}} \mathcal{L}_{\text{hm}} + \lambda_{\text{per}} \mathcal{L}_{\text{per}}, \quad (9)$$

where \mathcal{L}_{adv} denotes GAN loss. We set the $\lambda_{\text{id}} = 10$, $\lambda_{\text{lm}} = 10000$, $\lambda_{\text{hm}} = 100$, and $\lambda_{\text{per}} = 100$, respectively.

Experiments

Implementation details. The training images for Face-Controller are collected from CelebA-HQ (Karras et al. 2017), FFHQ (Karras, Laine, and Aila 2019), and VGGFace (Parkhi, Vedaldi, and Zisserman 2015) datasets. The face alignment for 3DMM coefficients extraction is processed according to (Deng et al. 2019b). The size of an aligned face is 224×224 . For region-wise style codes, we employ an off-the-shelf BiSeNet model (Yu et al. 2018) to obtain semantic labels, which provides 19 different region categories. The generator and discriminator are trained around 500K steps, respectively. More details can be found in the Supplementary Materials.

Qualitative Analysis

Face swapping. We first apply our method to the task of face swapping. Our method is compared with recent state-of-the-art methods including DeepFake (Petrov et al. 2020), FaceSwap¹, FSGAN (Nirkin, Keller, and Hassner 2019), and FaceShifter (Li et al. 2020). All methods are evaluated on the FaceForensics++ dataset (Rossler et al. 2019). As shown in Fig. 3, even though the rival methods are specifically designed for this problem, our method still obtains better or comparable results. Note that some methods, such as DeepFake and FSGAN, include a blending process to fuse the swapped faces with the backgrounds of target faces. Such a blending process is unstable and prone to failure, resulting in unpleasant illumination and artifacts. As contrast, our method and FaceShifter discard the blending process and achieve better results. Compared to FaceShifter, our method shows better preservation of identity because we have explicitly enhanced the disentanglement of attributes by carefully designing the representation and learning losses.

To demonstrate the power of our method on editing diverse face images, we also evaluate our method on face images downloaded from the Internet and faces generated by StyleGAN2 (Karras et al. 2020). The results are shown in Fig. 4. We observe that our method achieves very impressive results even when swapping faces with uncommon illumination or textures.

Region-wise style editing. The disentanglement of region-wise style codes enables our method to achieve high performance in makeup transfer. To validate this, we compare with recent methods such as LADN (Gu et al. 2019) and SEAN (Zhu et al. 2020b). All methods are evaluated on MT-dataset (Dantcheva, Chen, and Ross 2012). As illustrated in Fig. 5, our method appears to preserve more identity as compared to SEAN. Our method also ensures the fine-grained makeup transfer by utilizing region-wise style codes to adaptively adjust different face regions. This cannot be achieved by LADN, which uses a single entangled code to represent the whole face. To further demonstrate our capability of fine-grained makeup transfer, we show more results in Fig. 6. Here, we progressively transfer the styles of eyes, lips, and faces based on the results of the previous column. This can be simply implemented by editing corresponding region-wise style codes. These results also further verify the strong disentanglement of region-wise style codes.

Face manipulation and relighting. Our method can independently adjust the expression, pose, and illumination of faces. We demonstrate this ability by interpolating the above factors independently. The results are presented in Fig. 7. Based on the visual inspection of these results, we observe our method enables precise manipulation of the illumination and expression. To obtain a photorealistic background, we directly embed the background information into the generator rather than enforcing the network to learn how to inpaint. Therefore, The generated image may appear black holes along the edge of the face when the pose varies hugely.

¹<https://github.com/MarekKowalski/FaceSwap>




Figure 4: Face swapping on different sources of face images. Here, the first three images are generated by StyleGAN2 and others are collected from the Internet.




Figure 5: Comparison with other makeup transfer methods.

Quantitative Analysis

Face swapping. We report the quantitative results on FaceForensics++ dataset and compare our method with DeepFake, FaceSwap, FSGAN, and FaceShifter. We follow the metrics in FaceShifter. We uniformly sample 10 frames from each of 1000 videos and get 10K faces in total. Then, we evaluate the accuracy of identity retrieval (abbrev. Retr.) to check whether the identity of source face is preserved after face swapping; the expression (abbrev. Exp.) and pose error to check whether the swapped face keeps the expression and pose of the target face. We also use FID scores (Heusel et al. 2017) to evaluate the fidelity of swapped faces.

For identity retrieval, we apply CosFace (Wang et al. 2018) to extract identity embedding and retrieve the closest face by using cosine similarity. For pose and expression evaluation, we use a pose estimator to estimate head pose and an expression recognition model to extract expression embedding. Then, we report the L2 distance of pose vectors and expression embedding. The pose estimator is the same with that in FaceShifter. Since the expression model used in




Figure 6: Progressive makeup transfer.

FaceShifter is not open-source, we employ another model, i.e. (Meng et al. 2019), to enable expression error comparison. As shown in Tab. 1, our method achieves the best results in identity preservation and fidelity.

Region-wise style editing. We evaluate the performance of makeup transfer by users study. We mainly consider three dimensions: (1) the identity preservation; (2) the makeup transferability; (3) the fidelity of faces. All methods generate 1K images with the same source and reference faces. Five users with rich knowledge in image generation are asked to rate the best result in each case under each metric. Finally, the average scores are reported in Tab. 2. Since LADN is



Figure 7: The interpolated results with respect to illumination, expression, and poses from top to bottom.

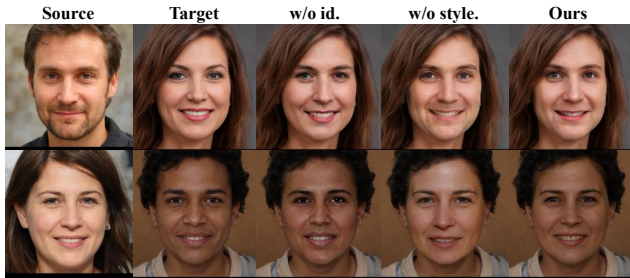


Figure 8: The face swapping results in ablation study. Here, w/o id. refers to our generator without identity encoder; w/o style. refers to our generator without style encoder.

specially designed to disentangle the identity and makeup, it enables the excellent preservation of identity and makeup. But compared with LADN, our method can generate higher-quality faces, as shown in Fig. 5. SEAN can perform very well on fine-grained makeup transfer. But it cannot precisely preserve the identity and often contains unnatural textures which affects its scores in all aspects.

Ablation Study

To verify the necessity of the identity encoder and style encoder, we conduct ablation experiments: the proposed generator with (1) no style encoder and (2) no identity encoder. Taking face swapping as an example, the results are demon-

Method	Retr.	Pose	Exp.	FID
DeepFake	81.96	4.14	0.47	4.29
FaceSwap	54.19	2.51	0.40	3.81
FSGAN	57.34	2.84	0.38	4.35
FaceShifter	97.38	2.96	0.36	4.05
Ours	98.27	2.65	0.39	3.51

Table 1: The accuracies of identity retrieval, pose and expression errors, and FID scores of different face swapping methods on FaceForensics++ dataset.

Method	ID	Style	Fidelity
SEAN	6.2	14.1	5.8
LADN	49.1	40.5	38.8
Ours	44.7	45.4	55.4

Table 2: Users study results w.r.t. identity preservation, makeup transfer, and fidelity.

Method	Retr.	Pose	Exp.	FID
w/o id.	25.97	2.57	0.36	4.12
w/o style.	98.21	2.83	0.41	4.84
Ours	98.27	2.65	0.39	3.51

Table 3: The accuracies of identity retrieval, pose and expression errors, and FID scores in ablation experiments.

strated in Fig. 8. If we discard the identity embedding, the generated images have identity gaps with source images, as shown in the third column of Fig. 8. Further, when we discard the style encoder, the generated styles aren't consistent with target image in regions like lips, as shown in the fourth column of Fig. 8. This indicates that 3DMM coefficients don't provide detailed texture for generating high-fidelity faces. The quantitative results are also consistent with the above conclusion. When style information isn't given, the identity is nearly not affected due to the disentangled representation. When identity information isn't provided, the accuracy of identity retrieval severely drops. Note that identity encoder and style encoder are irrelevant to pose and expression change, the pose and expression errors of these two models are very close to our full model. More ablation study results are shown in Appendix.

Limitations and Conclusion

Currently, our approach has some limitations that can be addressed in follow-up work. As we mentioned above, when changing pose hugely, the generated image may appear black holes. The gaze direction may be not consistent with the target image. This can be optimized by adding landmarks of the eyeball to constrain the gaze direction.

In conclusion, we propose a feed-forward network to generate photorealistic faces with desired attributes. This network designs a disentangled face representation and devises unsupervised losses to ensure the control of various face attributes. A variety of face applications can be implemented by this method.

Acknowledgments

This work was done when Zhiliang Xu was an intern at Baidu. This work was also supported by the National Program for Support of Top-notch Young Professionals and the Program for HUST Academic Frontier Youth Team 2017QYTD08, which are given to Dr. Xiang Bai.

References

- Abdal, R.; Qin, Y.; and Wonka, P. 2019. Image2StyleGAN: How to embed images into the StyleGAN latent space? In *Proc. CVPR*, 4432–4441.
- Blanz, V.; Vetter, T.; et al. 1999. A morphable model for the synthesis of 3D faces. In *Proc. SIGGRAPH*, volume 99, 187–194.
- Bouchacourt, D.; Tomioka, R.; and Nowozin, S. 2018. Multi-level variational autoencoder: Learning disentangled representations from grouped observations. In *Proc. AAAI*.
- Cao, C.; Weng, Y.; Zhou, S.; Tong, Y.; and Zhou, K. 2013. FaceWareHouse: A 3D facial expression database for visual computing. *TVCG* 20(3): 413–425.
- Chen, R. T.; Li, X.; Grosse, R. B.; and Duvenaud, D. K. 2018. Isolating sources of disentanglement in variational autoencoders. In *Proc. NeurIPS*, 2610–2620.
- Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *Proc. NeurIPS*, 2172–2180.
- Dantcheva, A.; Chen, C.; and Ross, A. 2012. Can facial cosmetics affect the matching accuracy of face recognition systems? In *Proc. BTAS*, 391–398. IEEE.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019a. Arcface: Additive angular margin loss for deep face recognition. In *Proc. CVPR*, 4690–4699.
- Deng, Y.; Yang, J.; Chen, D.; Wen, F.; and Tong, X. 2020. Disentangled and Controllable Face Image Generation via 3D Imitative-Contrastive Learning. In *Proc. CVPR*, 5154–5163.
- Deng, Y.; Yang, J.; Xu, S.; Chen, D.; Jia, Y.; and Tong, X. 2019b. Accurate 3D face reconstruction with weakly-supervised learning: From single image to image set. In *Proc. CVPR Workshops*, 0–0.
- EGGER, B.; SMITH, W. A.; TEWARI, A.; WUHRER, S.; ZOLLHOEFER, M.; BEELER, T.; BERNARD, F.; BOLKART, T.; KORTYLEWSKI, A.; ROMDHANI, S.; et al. 2020. 3D Morphable Face Models—Past, Present, and Future. *TOG* 39(5): 1–38.
- Gecer, B.; Bhattarai, B.; Kittler, J.; and Kim, T.-K. 2018. Semi-supervised adversarial learning to generate photorealistic face images of new identities from 3D morphable model. In *Proc. ECCV*, 217–234.
- Gu, Q.; Wang, G.; Chiu, M. T.; Tai, Y.-W.; and Tang, C.-K. 2019. LADN: Local adversarial disentangling network for facial makeup and de-makeup. In *Proc. ICCV*, 10481–10490.
- Ha, S.; Kersner, M.; Kim, B.; Seo, S.; and Kim, D. 2019. MarioNETte: Few-shot face reenactment preserving identity of unseen targets. In *Proc. AAAI*.
- He, Z.; Zuo, W.; Kan, M.; Shan, S.; and Chen, X. 2019. AttGAN: Facial attribute editing by only changing what you want. *TIP* 28(11): 5464–5478.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Proc. NeurIPS*, 6626–6637.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; et al. 2015. Spatial transformer networks. In *Proc. NeurIPS*, 2017–2025.
- Johnson, J.; Alahi, A.; and Fei-Fei, L. 2016. Perceptual losses for real-time style transfer and super-resolution. In *Proc. ECCV*, 694–711. Springer.
- 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*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proc. CVPR*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 8110–8119.
- Kim, H.; Garrido, P.; Tewari, A.; Xu, W.; Thies, J.; Niessner, M.; Pérez, P.; Richardt, C.; Zollhöfer, M.; and Theobalt, C. 2018. Deep video portraits. *TOG* 37(4): 1–14.
- Klys, J.; Snell, J.; and Zemel, R. 2018. Learning latent subspaces in variational autoencoders. In *Proc. NeurIPS*, 6444–6454.
- Li, L.; Bao, J.; Yang, H.; Chen, D.; and Wen, F. 2020. FaceShifter: Towards high fidelity and occlusion aware face swapping. In *Proc. CVPR*.
- Lin, Z.; Thekumparampil, K. K.; Fanti, G.; and Oh, S. 2019. InfoGAN-CR: Disentangling Generative Adversarial Networks with Contrastive Regularizers. *arXiv preprint arXiv:1906.06034*.
- Locatello, F.; Bauer, S.; Lucic, M.; Raetsch, G.; Gelly, S.; Schölkopf, B.; and Bachem, O. 2019. Challenging common assumptions in the unsupervised learning of disentangled representations. In *Proc. ICML*, 4114–4124.
- Meng, D.; Peng, X.; Wang, K.; and Qiao, Y. 2019. Frame attention networks for facial expression recognition in videos. In *Proc. ICIP*, 3866–3870. IEEE.
- Nguyen-Phuoc, T.; Li, C.; Theis, L.; Richardt, C.; and Yang, Y.-L. 2019. HoloGAN: Unsupervised learning of 3D representations from natural images. In *Proc. ICCV*, 7588–7597.
- Nie, W.; Karras, T.; Garg, A.; Debhath, S.; Patney, A.; Patel, A. B.; and Anandkumar, A. 2020. Semi-Supervised StyleGAN for Disentanglement Learning. *arXiv arXiv:2003*.
- Nirkin, Y.; Keller, Y.; and Hassner, T. 2019. FSGAN: Subject Agnostic Face Swapping and Reenactment. In *Proc. ICCV*, 7184–7193.

- Park, T.; Liu, M.; Wang, T.; and Zhu, J. 2019. Semantic Image Synthesis With Spatially-Adaptive Normalization. In *Proc. CVPR*, 2337–2346.
- Parkhi, O. M.; Vedaldi, A.; and Zisserman, A. 2015. Deep face recognition. In *Proc. BMVC*.
- Petrov, I.; Gao, D.; Chervoniy, N.; Liu, K.; Marangonda, S.; Umé, C.; Jiang, J.; RP, L.; Zhang, S.; Wu, P.; et al. 2020. DeepFaceLab: A simple, flexible and extensible face swapping framework. *arXiv preprint arXiv:2005.05535*.
- Razavi, A.; van den Oord, A.; and Vinyals, O. 2019. Generating diverse high-fidelity images with VQ-VAE-2. In *Proc. NeurIPS*, 14866–14876.
- Risser, E.; Wilmot, P.; and Barnes, C. 2017. Stable and controllable neural texture synthesis and style transfer using histogram losses. *arXiv preprint arXiv:1701.08893*.
- Rossler, A.; Cozzolino, D.; Verdoliva, L.; Riess, C.; Thies, J.; and Nießner, M. 2019. FaceForensics++: Learning to detect manipulated facial images. In *Proc. ICCV*, 1–11.
- Shen, W.; and Liu, R. 2017. Learning residual images for face attribute manipulation. In *Proc. CVPR*, 4030–4038.
- Shen, Y.; Gu, J.; Tang, X.; and Zhou, B. 2020. Interpreting the latent space of GANs for semantic face editing. In *Proc. CVPR*, 9243–9252.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sun, R.; Huang, C.; Shi, J.; and Ma, L. 2018. Mask-aware photorealistic face attribute manipulation. *arXiv preprint arXiv:1804.08882*.
- Tewari, A.; Elgharib, M.; Bharaj, G.; Bernard, F.; Seidel, H.-P.; Pérez, P.; Zollhofer, M.; and Theobalt, C. 2020. StyleRig: Rigging StyleGAN for 3D Control over Portrait Images. In *Proc. CVPR*, 6142–6151.
- Thies, J.; Zollhofer, M.; Stamminger, M.; Theobalt, C.; and Nießner, M. 2016. Face2Face: Real-time face capture and reenactment of RGB videos. In *Proc. CVPR*, 2387–2395.
- Wang, H.; Wang, Y.; Zhou, Z.; Ji, X.; Gong, D.; Zhou, J.; Li, Z.; and Liu, W. 2018. Cosface: Large margin cosine loss for deep face recognition. In *Proc. CVPR*, 5265–5274.
- Xu, S.; Yang, J.; Chen, D.; Wen, F.; Deng, Y.; Jia, Y.; and Tong, X. 2020. Deep 3D Portrait from a Single Image. In *Proc. CVPR*, 7710–7720.
- Yu, C.; Wang, J.; Peng, C.; Gao, C.; Yu, G.; and Sang, N. 2018. BiSeNet: Bilateral segmentation network for real-time semantic segmentation. In *Proc. ECCV*, 325–341.
- Zeng, X.; Pan, Y.; Wang, M.; Zhang, J.; and Liu, Y. 2020. Realistic Face Reenactment via Self-Supervised Disentangling of Identity and Pose. In *Proc. AAAI*.
- Zhang, J.; Huang, Y.; Li, Y.; Zhao, W.; and Zhang, L. 2019. Multi-attribute transfer via disentangled representation. In *Proc. AAAI*, volume 33, 9195–9202.
- Zhu, J.; Shen, Y.; Zhao, D.; and Zhou, B. 2020a. In-domain GAN inversion for real image editing. In *Proc. ECCV*.
- Zhu, P.; Abdal, R.; Qin, Y.; and Wonka, P. 2020b. SEAN: Image Synthesis with Semantic Region-Adaptive Normalization. In *Proc. CVPR*, 5104–5113.
- Zhu, Z.; Xu, Z.; You, A.; and Bai, X. 2020c. Semantically Multi-modal Image Synthesis. In *Proc. CVPR*, 5467–5476.