

BlendFace: Re-designing Identity Encoders for Face-Swapping

Kaede Shiohara^{1*} Xingchao Yang² Takafumi Taketomi²
¹The University of Tokyo ²CyberAgent AI Lab

shiohara@cvm.t.u-tokyo.ac.jp {you.koutyo,taketomi_takafumi}@cyberagent.co.jp

Abstract

The great advancements of generative adversarial networks and face recognition models in computer vision have made it possible to swap identities on images from single sources. Although a lot of studies seems to have proposed almost satisfactory solutions, we notice previous methods still suffer from an identity-attribute entanglement that causes undesired attributes swapping because widely used identity encoders, e.g., ArcFace, have some crucial attribute biases owing to their pretraining on face recognition tasks. To address this issue, we design BlendFace, a novel identity encoder for face-swapping. The key idea behind BlendFace is training face recognition models on blended images whose attributes are replaced with those of another mitigates inter-personal biases such as hairstyles. BlendFace feeds disentangled identity features into generators and guides generators properly as an identity loss function. Extensive experiments demonstrate that BlendFace improves the identity-attribute disentanglement in face-swapping models, maintaining a comparable quantitative performance to previous methods. The code and models are available at <https://github.com/mapoon/BlendFace>.

1. Introduction

Face-swapping aims to replace target identities with source identities in images while preserving the target attributes, e.g., facial expression, hair, pose, and background. This task is receiving considerable attention because of its potential applications in various fields, such as films and metaverses. Recent advances in generative adversarial networks (GANs) [7, 25, 34, 36, 51, 58, 89] have enabled the photo-realistic image generation in various conditions, e.g., attribute [29], identity [9], and expression [83], as well as unconditional image generation. Moreover, the advancement in face recognition models provides powerful identity encoders for face-swapping, which boosts the trans-



Figure 1: **Examples of attribute leakages.** The direct use of face recognition models, i.e., ArcFace [17] as an identity guidance causes attribute leakages, especially on head shapes and hairstyles due to attribute biases inherent in the face recognition models whereas our identity encoder BlendFace address these problems. Best viewed in zoom.

ferability of identities from source inputs to generated images and leads to successful one-shot face-swapping models [9, 14, 22, 45, 50, 72, 74–76, 90] with reasonable quality.

However, despite these impressive efforts, a critical issue still remains. Previous state-of-the-art methods suffer from identity-attribute entanglements because of biased guidance from face recognition models used as identity encoders. Fig. 1 presents the failure cases of a traditional face recognition model ArcFace [17]. As shown in the figure, ArcFace-based face-swapping models tend to swap undesired attributes, e.g., hairstyles and head shapes. This is because images of the same identity have strong correlations for some attributes; therefore, face recognition models accidentally learn to recognize the attributes as identities, which causes misguidances in training face-swapping models. Though certain studies in the field of face recognition propose effective approaches to mitigate biases between individuals, they cannot be solutions for biases in face-swapping models as they do not consider intra-personal biases.

In this paper, we propose BlendFace, a novel identity encoder that provides well-disentangled identity features for

*Work done during an internship at CyberAgent AI Lab

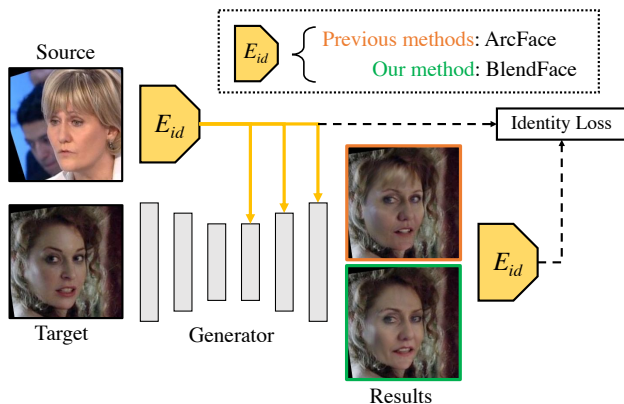


Figure 2: **Overview of face-swapping models.** Identity encoder E_{id} is used to extract identity features from source images and to guide generators as an identity distance loss. We replace ArcFace with our identity encoder BlendFace to achieve more disentangled face-swapping.

face-swapping. First, we analyze the widely used identity encoder ArcFace [17] on VGGFace2 [12] dataset. The comparison of identity similarity distributions using pseudo-swapped images clarifies that attributes such as hairstyles, colors, and head shapes strongly affect their similarities because of attribute biases in ArcFace, which is expected to prevent face-swapping models from disentangling identity and attribute. Based on the observation from the preliminary experiment, we design BlendFace by simply training ArcFace with swapped images so that the model does not focus on attributes of faces, which bridges the gap between similarity distributions of swapped faces and real faces. We then train a face-swapping model using BlendFace that performs as a source feature extractor and identity guidance in the loss function. As shown in Fig. 2, by replacing the traditional identity encoder in the source feature extraction and loss computation with BlendFace, face-swapping models are trained to generate more disentangled face-swapping results. Importantly, our work is compatible with previous face-swapping studies; BlendFace can be applied to various learning-based face-swapping models.

In the experiment part, we compare our model with state-of-the-art face-swapping models on FaceForensics++ [60] dataset following the convention. The comparison demonstrates that the proposed method is superior to or on par with the competitors in identity similarity and attribute preservation, *i.e.*, expression, pose, and gaze, while improving visual consistency of swapped results compared to previous models. In addition, our ablation study and analysis prove the advantages of BlendFace from various perspectives for face-related research.

2. Related Work

Face-Swapping. Face manipulations, particularly face-swapping, is an area of significant research in computer vision owing to its potential applications such as realistic digital avatar creation. Early methods employ traditional image processing [10] and 3D morphable models (3DMMs) [11, 15, 55]. The brilliant successes of generative adversarial networks (GANs) [7, 25, 34, 36, 51, 58, 89] in computer vision has driven the extensive exploration of learning-based face-swapping models. FSGAN [54] realizes subject-agnostic face-swapping via four encoder-decoder networks, *i.e.*, reenactment, segmentation, inpainting, and blending networks. Sophisticated face recognition models [17, 41, 49, 68] that learn rich identity information from large-scale facial datasets [6, 12, 26, 79, 91] improves the identity preservation in face swapping. Sim-Swap [14] proposes a weak feature matching loss between each generated image and target image in the discriminator’s feature space to balance the preservation of the source identity and target attribute. FaceShifter [45] proposes a two-stage framework including AEI-Net that blends the features extracted from source and target images in multiple scales and HEAR-Net that learns to reconstruct occlusions using objects datasets [8, 13, 21]. InfoSwap [22] introduces information-theoretical loss functions to disentangle identities. HifiFace [72] incorporates a 3DMM [18] to its identity extraction to retrain the source appearance and shape. Smooth-Swap [40] develops a smooth identity encoder to stably GAN-training using self-supervised pre-training [39]. Recently, some studies [74, 75, 90] reveal pre-trained StyleGANs [35–37] provide strong priors to generate photo-realistic facial images at megapixel resolution for face-swapping. MegaFS [90] generates swapped faces by replacing high semantic features of target images with those of source images. RAFSwap [74] integrates semantic-level features with a face-parsing model [80]. FSLSD [75] transfers multi-level attributes via side-outputs from StyleGAN. StyleSwap [76] proposes iterative identity optimization that effectively preserves source identities. StyleFace [50] and UniFace [73] unify face-swapping into de-identification and reenactment, respectively. In this paper, we re-design identity encoders independently of these state-of-the-art approaches; our encoder can be easily incorporated into previous learning-based face-swapping models.

Face Recognition. The task of face recognition is a fundamental problem in the research field. Recent approaches have mainly been conducted using deep convolutional networks. In particular, margin-based loss functions (*e.g.*, [17, 49, 68]) significantly enhance the performance of face recognition. However, some studies [33, 59, 63, 64, 67, 82] have found such identity encoders contain biases of attributes, *e.g.*, pose, hairstyles, color, races,

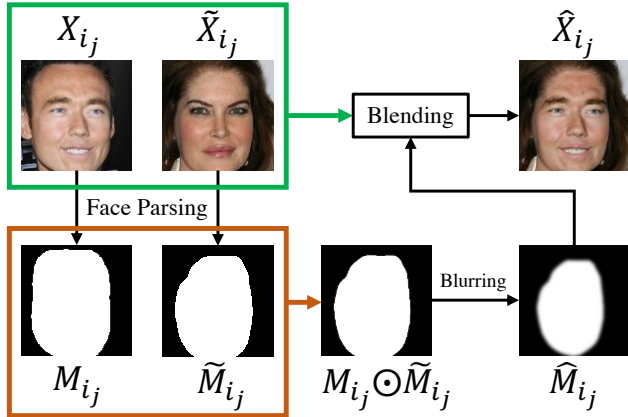


Figure 3: **The pipeline for pseudo-positive samples.** We first generate masks M_{ij} and \tilde{M}_{ij} corresponding to X_{ij} and \tilde{X}_{ij} , respectively, then generate a smooth mask \hat{M}_{ij} by blurring the intersection mask $M_{ij} \odot \tilde{M}_{ij}$. Finally we blend X_{ij} and \tilde{X}_{ij} with \hat{M}_{ij} to generate the blended image \hat{X}_{ij} .

and gender; therefore debiasing face recognition models have been a concern topic in the field. IMAN [70] proposes information-theoretical adaptation networks. RL-RBN [69] adapts the margin of ArcFace by reinforcement learning. GAC [24] proposes adaptive layers comprising demographic-group-specific kernels and attention modules. DebFace [23] disentangles gender, age, race, and identity using feature disentangling blocks and aggregation blocks in an adversarial learning.

Although these methods effectively mitigate the biases between identities, they however do not focus on interpersonal biases. Therefore, existing identity extractors cause undesired attributes swapping because images of each identity in datasets used for face recognition have strong correlations in some attributes, *e.g.*, hairstyles, colors, and head shapes. To solve this problem, we design an debiased encoder that extracts disentangled identity features from facial images by training a face recognition model with synthetic images that have swapped attributes, which enables well-disentangled face-swapping.

3. Attribute Bias in Face Recognition Models

Given source and target images, face-swapping aims to generate a facial image where the target identity is replaced with the source identity while preserving the attributes of the target image. First of all, we rethink the identity encoding for face-swapping by conducting a preliminary experiment with ArcFace [17] adopted by most face-swapping models, *e.g.*, [14, 22, 45, 72, 75, 90]. The key observation is replacing attributes of one individual with those of another causes a degradation of the identity similarity, which indi-

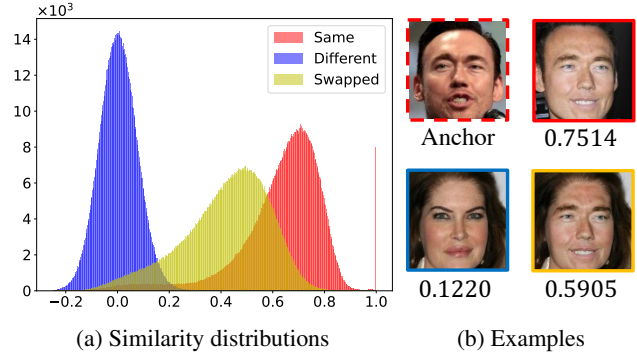


Figure 4: **Analysis of ArcFace on VGGFace2.** ArcFace tends to underestimate the identity similarities between anchor images and swapped faces.

cates attribute biases inherent in ArcFace.

3.1. Identity Distance Loss

One of the difficulties in face-swapping is the absence of ground truth images. Given two images of different identities for source and target inputs during the training, generated images are constrained by some feature-based losses, *e.g.*, appearance [17, 32, 85], 3D face shapes [18], and segmentation [80] to preserve identities of source and attributes of target. Notably, most of the previous methods adopt ArcFace [17] trained on large-scale face recognition datasets [6, 12, 26, 79, 91] to extract identity information from source inputs and measure the identity distance between a source image X_s and swapped image $Y_{s,t}$ as follows:

$$\mathcal{L}_{id} = 1 - \cos\langle E_{id}(X_s), E_{id}(Y_{s,t}) \rangle, \quad (1)$$

where E_{id} denotes ArcFace encoder and $\cos\langle u, v \rangle$ is the cosine similarity of vectors u and v .

3.2. Analysis of Identity Similarity

Here, we explore the attribute biases of ArcFace [17] on VGGFace2 [12] dataset from a perspective of face-swapping. As shown in Fig. 3, we first randomly sample the j -th image of the i -th identity which is presented as X_{ij} , and then compute cosine similarities between X_{ij} and all the images of the same identity $\{X_{i_1}, X_{i_2}, \dots, X_{i_{n_i}}\}$, where n_i denotes the number of images of the identity i . Subsequently, inspired by Face X-ray [46], we search for an image \tilde{X}_{ij} with the closest facial landmarks to X_{ij} from randomly sampled 100 images whose identities are not i for each X_{ij} . After transferring the color statistics μ and σ in the Lab space of X_{ij} to that of \tilde{X}_{ij} , we replace the face of \tilde{X}_{ij} with that of X_{ij} by blending them with a mask \hat{M}_{ij} generated by multiplying the inner mask M_{ij} of X_{ij} by \tilde{M}_{ij}

of \tilde{X}_{i_j} :

$$\hat{X}_{i_j} = X_{i_j} \odot \hat{M}_{i_j} + \tilde{X}_{i_j} \odot (1 - \hat{M}_{i_j}), \quad (2)$$

where \odot denotes the point-wise product, \tilde{X}_{i_j} represents the synthetic swapped image, and $\hat{M}_{i_j} = \text{Blur}(M_{i_j} \odot \tilde{M}_{i_j})$. We then compute the cosine similarities between X_{i_j} and replaced images $\{\hat{X}_{i_1}, \hat{X}_{i_2}, \dots, \hat{X}_{i_{n_i}}\}$. In addition, we calculate cosine similarities between \tilde{X}_{i_j} and the closest images \tilde{X}_{i_j} . We repeat this procedure for all identities and plot the similarity distributions in Fig. 4a.

The key observations of this experiment are as follow: 1) The similarities are equal to one only for identical image comparisons; otherwise, the similarities are almost lower than 0.85 even if two images have the same identity. 2) The similarities between anchor images and synthetic positive images are lower than those between actual positive pairs, indicating that the color distribution of faces and attributes of outer face region strongly affect the similarities. This is because face recognition models tend to recognize certain attributes, *e.g.*, hairstyles and head shapes as identities because an image set of each identity of face recognition datasets used for training often has correlations in their attributes.

These results leads us to assume that minimizing the identity loss in Eq. 1 using traditional face recognition models, *e.g.*, ArcFace [17] conflicts with the presevation of target attributes because face recognition models have attribute biases, which constrains generators excessively to transfer not only identities but also attributes from source images.

4. BlendFace

We propose a novel identity encoder BlendFace to solve the problem of the identity-attribute entanglement owing to the attribute biases in face recognition models as discussed in Sec. 3.2. First, we introduce a pre-training strategy to train debiased identity encoders using synthetically swapped faces. We then incorporate our identity encoder into a face-swapping model disentangling attributes and identities for high-fidelity face-swapping.

4.1. Pre-training with Swapped Faces

As discussed in Sec. 3.2, traditional face recognition models trained on real face datasets, *e.g.*, MS-Celeb1M [26] learn accidentally attribute biases because images of each identity are highly correlated with each other in some attributes, *e.g.*, hairstyles and makeup; this produces poor results for source and target images with large attribute differences. To tackle this problem, we develop a debiased identity encoder BlendFace that can be achieved by training a face recognition model with synthetic facial images whose attributes are swapped. We adopt ArcFace [17] as

our base model and train it with blended images that have synthetically swapped attributes. For each sample during the training, we swap attributes of input images in the same manner as in Sec. 3.2 with probability p . The loss function of ArcFace [17] is as follows:

$$\mathcal{L} = -\log \frac{e^{s \cos(\theta_{y_i} + m)}}{e^{s \cos(\theta_{y_i} + m)} + \sum_{k=1, k \neq y_i}^K e^{s \cos \theta_k}}, \quad (3)$$

where θ_{y_i} represents the angle between the deep feature vector and weight vector of the encoder. K , s , and m denote the number of classes, scale, and margin, respectively. We observe that our pretraining bridges the gap between the distribution of ‘‘Swapped’’ and that of ‘‘Same’’ (see Fig. 7b). We conduct ablations for p and \hat{M}_{i_j} in Sec. 5.4.

4.2. Face-Swapping with BlendFace

To validate the effectiveness of BlendFace, we construct a face-swapping model with BlendFace. We denote source, target, and generated images as X_s , X_t , and $Y_{s,t}$ ($= G(X_s, X_t)$), respectively. We adopt a state-of-the-art architecture AEI-Net [45] with some modifications. We replace ArcFace used in encoding source identities and computing the distance loss \mathcal{L}_{id} (Eq. 1) with BlendFace. We incorporate a blending mask predictor into the attributes encoder inspired by previous studies (*e.g.*, [72, 74, 76]). A predicted mask \hat{M} is supervised by the binary cross entropy loss \mathcal{L}_{mask} with the ground truth mask M from a face-parsing model [80] as follows:

$$\mathcal{L}_{mask} = -\sum_{x,y} \{M_{x,y} \log \hat{M}_{x,y} + (1 - M_{x,y}) \log (1 - \hat{M}_{x,y})\}, \quad (4)$$

where x and y are the spacial coordinates of image. We feed different images that share the same identity for source and target inputs rather than feeding the same image when activating the reconstruction loss as follows:

$$\mathcal{L}_{rec} = \begin{cases} \|X_t - Y_{s,t}\|_1 & \text{if } ID(X_t) = ID(X_s), \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

We sample the same identity for source and target images with $p = 0.2$. We use the cycle consistency loss instead of the attributes loss used in the original FaceShifter as follows:

$$\mathcal{L}_{cyc} = \|X_t - G(X_t, Y_{s,t})\|_1. \quad (6)$$

We use the same adversarial loss term \mathcal{L}_{adv} as in GauGAN [56]. The total loss \mathcal{L} is formulated as:

$$\mathcal{L} = \mathcal{L}_{adv} + \mathcal{L}_{mask} + \lambda_1 \mathcal{L}_{id} + \lambda_2 \mathcal{L}_{rec} + \lambda_3 \mathcal{L}_{cyc}, \quad (7)$$

where the coefficients λ_1 , λ_2 , and λ_3 are hyper-parameters that balance the loss functions. Please see the supplementary material for a more detailed description of the architecture of our face-swapping model.

Model	Identity Distance						Attribute Distance					
	Arc	Arc-R	Blend	Blend-R	Shape	Shape-R	Expr	Expr-R	Pose	Pose-R	Gaze	Gaze-R
Deepfakes [1]*	<u>0.5388</u>	<u>0.4048</u>	0.3704	0.2859	0.3339	0.4327	0.2037	-	0.0374	-	0.2891	-
FaceSwap [2]*	0.5916	0.4594	0.4653	0.3606	0.3355	0.4546	0.1824	-	0.0189	-	0.2273	-
FSGAN [54]	0.7188	0.5043	0.6640	0.4464	0.3918	0.5410	0.1419	0.4212	0.0173	0.1835	0.1772	0.4120
FaceShifter [45]*	0.3826	0.3197	0.3143	0.2607	0.3190	0.4516	0.1720	-	0.0162	-	0.1840	-
SimSwap [14]	0.3594	0.3012	0.3710	0.2992	0.3309	0.4761	0.1582	0.4809	<u>0.0139</u>	<u>0.1512</u>	<u>0.1599</u>	<u>0.4052</u>
HifiFace [72]*	0.3593	0.2940	<u>0.3488</u>	<u>0.2771</u>	<u>0.3119</u>	0.4405	0.1730	0.5066	0.0161	0.1702	0.1663	0.4183
MegaFS [90]	0.5905	0.3802	0.6112	0.3869	0.3499	<u>0.4221</u>	0.2057	0.5408	0.0377	0.2455	0.1607	0.4066
InfoSwap [22]	0.4308	0.3056	0.4708	0.3223	0.2962	0.3963	0.2007	0.5719	0.0192	0.1978	0.1748	0.4250
FSLSD [75]	0.7952	0.4954	0.7853	0.4623	0.3706	0.4639	0.2023	0.5385	0.0231	0.2407	0.2598	0.5405
Ours	0.4630	0.3987	0.2574	0.2316	0.3542	0.5107	<u>0.1537</u>	<u>0.4665</u>	0.0125	0.1395	0.1301	0.3514

Table 1: **Comparison with state-of-the-art methods.** **Bold** and underlined values correspond to the best and the second-best values, respectively. Gray values are excluded from the evaluation because of the use of the same encoders in their training. * denotes officially released generated videos. Our method outperforms previous state-of-the-arts in pose and gaze, and achieves the second best results in expression.

5. Experiment

We validate the effectiveness of our method through extensive comparisons with previous methods, ablations and analyses. The results demonstrate that BlendFace improves the fidelity of identity similarity and attribute preservation compared with previous models.

5.1. Implementation Detail

Pretraining of BlendFace. We adopt MS-Celeb-1M [26] dataset to train BlendFace. The batch size and the number of epochs are set to 1024 and 20, respectively. We train our encoder on the loss in Eq. 3 for 20 epochs using SGD optimizer with learning rate 0.1. We set the probability p of replacing attributes to 0.5.

Training of face-swapping model. We adopt VG-GFace2 [12] dataset to train our face-swapping model. We align and crop the images following FFHQ [36] preprocessing. We use ADAM [42] optimizer with $\beta_1 = 0$, $\beta_2 = 0.999$, and $lr = 0.0004$ for our generator and discriminator. We train our model for 300k iterations with a batch size of 32. The coefficients of the total loss in Eq. 7 are empirically set to $\lambda_1 = 10$, $\lambda_2 = 5$, and $\lambda_3 = 5$.

5.2. Experiment Setup

Setup. Following the conventional evaluation protocol, we evaluate face-swapping models on FaceForensics++ (FF++) [60] dataset which includes 1000 real videos and 1000 generated videos each of Deepfakes [1], Face2Face [66], FaceSwap [2], NeuralTextures [65], and FaceShifter [45]. We follow the setting of pairs of source and target defined by the original FF++ dataset. And we use the same source frames as in HifiFace [72].

Metric. To evaluate the fidelity of generated images,

we consider six metrics: ArcFace (Arc) [17], BlendFace (Blend), face shape [18], expression (Expr) [18], head pose [18], and gaze [3]. We measure the distances between source and swapped images for identity metrics, *i.e.* ArcFace, BlendFace, and shape, and between target and swapped images for attribute metrics, *i.e.*, expression, pose, and gaze. We calculate the cosine distances of extracted feature vectors for ArcFace and BlendFace, L1 distances of predicted 3DMM parameters for shape, expression, and pose, and L1 distances of predicted Euler angles for gaze. We further compute the relative distances [40] that consider both source and target in all the metrics, which are denoted as “-R”.

5.3. Comparison with Previous Methods

Baselines. We compare our method with publicly available state-of-the-art models, *i.e.*, FSGAN [54], SimSwap [14], MegaFS [90], InfoSwap [22], and FSLSD [75]. We also adopt generated videos of Deepfakes [1], FaceSwap [2], and FaceShifter [45] from FF++ dataset and HifiFace [72] from the official project page [4].

Result. The quantitative result is presented in Table 1. Our model outperforms previous methods in terms of absolute and relative metrics of pose and gaze, and achieves the second-best result in absolute and relative metrics of expression. Although it is difficult to compare our model with previous methods in identity similarity using face recognition models because evaluations using the same encoder as in training tends to be overestimated, our model outperforms Deepfakes [1], FaceSwap [2], and FSGAN [54] in ArcFace at least. We also show some generated images in Fig. 5. We omit examples of Deepfakes, FaceSwap, and FaceShifter because these results are generated by different source frames. We can see that previous approaches suffer




Figure 5: **Qualitative comparison on FF++**. Our model performs consistent face-swapping while preserving target attributes, *e.g.*, hairstyles, expression, pose, and gaze. More results are included in the supplementary material.

from attribute leakages, particularly in changing hairstyles, (*e.g.*, (a) and (d)) and head shapes, (*e.g.*, (c) and (e)), whereas our method succeeds in generating robust results qualitatively. In addition, we observe that the StyleGAN2-based methods, *i.e.*, MegaFS [90] and FSLSD [75] are vulnerable to unseen domain because the models sometimes fail in the inversion of source identities. We include more generated results in the supplementary material due to the space limitations. Because our method limits face shape changes to maintain consistency between the inside and outside of faces, the similarities of face shapes between generated images and source images are slightly lower than those of other methods, which is one of the limitations of our approach discussed in Sec. 6.

5.4. Ablation and Analysis

Choices of source encoder and loss. We found that the performance of face-swapping models strongly rely on choices of identity encoders to extract source features and compute the identity loss. We train four face-swapping models with ArcFace and BlendFace set to source encoder or loss. Then

Setting		Distance			
Source	Loss	Arc	Blend	Pose	Gaze
ArcFace	ArcFace	0.3314	0.2918	0.0139	0.1450
BlendFace	ArcFace	0.4254	0.2730	<u>0.0128</u>	<u>0.1376</u>
ArcFace	BlendFace	<u>0.4123</u>	<u>0.2700</u>	0.0131	0.1377
BlendFace	BlendFace	0.4630	0.2574	0.0125	0.1301

Table 2: **Choices of source encoder and loss.** Our method achieves the best results in distances of BlendFace, pose, and gaze.

we evaluate these models in the same protocol as Sec. 5.3. Note that our model in Sec. 5.3 sets BlendFace both to identity encoder and loss. The result is given in Table 2. It can be observed that using ArcFace both in the source encoder and loss computation brings the worst results in distances of BlendFace, expression, pose, and gaze, which implies the identity-attribute entanglement of ArcFace. Our method meanwhile achieves the best preservation in pose and gaze though worst identity distance on ArcFace because ArcFace



Figure 6: **Choices of encoder and loss.** The notation “A-B” indicates using encoders A and B for source feature extraction and loss computation, respectively. Our model produces reasonable results in the identity similarity and attribute preservation. Best viewed in zoom.

tends to underestimate well-disentangled swapped results. We also show the generated images in Fig. 6. We can observe computing the identity loss with ArcFace produces inconsistent blending boundaries and attributes because of the entanglement of identities and attributes, while BlendFace produces consistent results. These results clearly indicate that our BlendFace outperforms ArcFace quantitatively and qualitatively in face-swapping.

Probability of replacing attributes. We examine the effect of the probability p of replacing attributes in pretraining of the identity encoder. We additionally train our encoders with $p = 0.25, 0.75,$ and $1.00,$ then we conduct the same experiment as Sec. 3.2 and plot the distributions in Fig 7. It can be observed that training identity encoders with swapped faces bridges the gap between the similarity distributions of actual positive samples denoted as “Same” and swapped faces denoted as “Swapped” in the figure, even when $p = 0.25$. Also, increasing p brings these distributions closer. We again emphasize that ArcFace [17], which corresponds to the case of $p = 0.0,$ underestimates swapped faces as shown in Fig. 4. This result supports that our pre-training approach removes attribute biases from face recognition models.

Face verification on real face datasets. We validate our encoders in the task of face verification. We adopt well-known benchmarks including LFW [30], CFP-FP [61], and AgeDB [53]. We compare our BlendFace with ArcFace




Figure 7: **Effect of the ratio of synthetic faces.** Training identity encoders with swapped faces brings the distributions of swapped faces and that of real positive faces closer together.

Model	LFW	CFP-FP	AgeDB
ArcFace [17]	0.9978	<u>0.9900</u>	0.9838
SynFace [57]	0.9763	0.8731	0.8243
BlendFace ($p = 0.25$)	<u>0.9983</u>	0.9904	0.9815
BlendFace ($p = 0.50$)	0.9984	0.9897	0.9813
BlendFace ($p = 0.75$)	0.9983	0.9891	<u>0.9828</u>
BlendFace ($p = 1.00$)	0.9977	0.9886	0.9800

Table 3: **Face verification result.** Our encoders keep the ability of face verification on real face datasets and outperforms SynFace, which support our encoders perform properly as identity guidance in training face-swapping models.

and SynFace [57], a variant of ArcFace trained on GAN-synthesized images. We adopt the *unrestricted with labelled outside data* protocol to evaluate models, following the convention in the research field of face recognition (*e.g.*, [17]). We give the results in Table 3. We observe that our encoders retain the ability of face verification despite the slight performance degradation. This is because the removed features, *e.g.*, for hairstyles, face-shapes, and colors, are useful in verifying real faces.

Blending mask. We examine the effect of blending masks in pretraining of BlendFace. As described in Sec. 3.2, we use a mask \tilde{M}_{i_j} generated from the intersection $M_{i_j} \odot \tilde{M}_{i_j}$ to blend source image X_{i_j} and target image \tilde{X}_{i_j} during pre-training. We here train BlendFace with M_{i_j} or \tilde{M}_{i_j} instead



Figure 8: **Effect of blending masks.** Using intersection mask $M_{i_j} \odot \tilde{M}_{i_j}$ contributes to the consistency of generated images.

of $M_{i_j} \odot \tilde{M}_{i_j}$. Then we train face-swapping models with these encoders. As shown in Fig. 8, we found the model with $M_{i_j} \odot \tilde{M}_{i_j}$ produces more consistent results than M_{i_j} or \tilde{M}_{i_j} . This is because blending X_{i_j} and \tilde{X}_{i_j} with M_{i_j} or \tilde{M}_{i_j} produces artifacts in the blended image \hat{X}_{i_j} during pretraining, which harms face-swapping models.

Saliency Map. To explore the effectiveness of BlendFace, we visualize the saliency maps of ArcFace and BlendFace. Inspired by sliding occlusions [84], we measure the averaged sensitivity of identity similarities between the occluded input images and the reference image over multiple mask sizes $\{16, 24, 32, 40, 48, 56\}$. As shown in Fig. 9, ArcFace focuses on both the inner and outer faces while our encoder properly does only on inner faces. The result supports that our BlendFace can swap only inner faces without the undesired attribute transfer of outer faces.

6. Limitations

Our novel identity encoder BlendFace provides disentangled identity features that are beneficial face-swapping and other face-related tasks; however, we notice some limitations of our method. First, our model can hardly change face shapes because we limit the region where identities are swapped to improve spatial consistencies between inside and outside of faces. Therefore when source and target images that have to different face shapes are input into our model, the generated image looks like the source subject in terms of the appearance but may not in terms of the face shape. Second, similar to previous methods, our method sometimes fail to preserve hard occlusions such as hands because of the lack of training samples of extreme scenes. It can be improved by incorporating HEAR-Net [45] into our model.

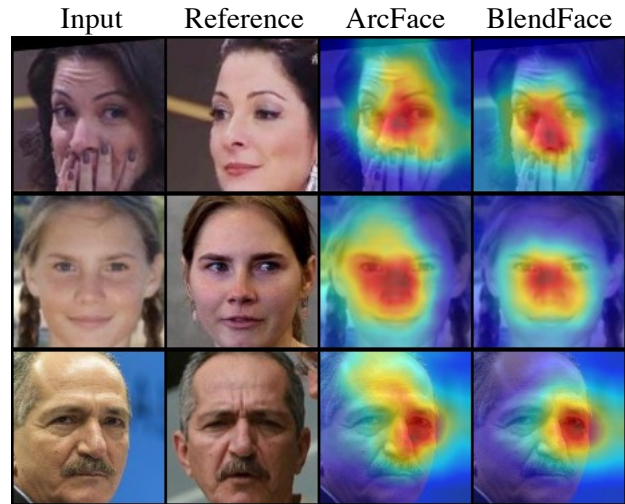


Figure 9: **Saliency maps of ArcFace and BlendFace.** ArcFace pays attention not only to inner faces but also to outer faces while BlendFace does only to inner faces.

7. Conclusion

In this paper, we present BlendFace, a well-disentangled identity encoder for more consistent face-swapping. The key observation behind BlendFace is that traditional face recognition models trained on real face datasets have biases in some attributes, *e.g.*, hairstyles and face-shapes, which leads inconsistent results in face-swapping. To tackle this problem, we train a face recognition model with blended images as pseudo-positive samples that have swapped attributes so that the encoder focuses only on inner faces, which improves the disentanglement of identity and attribute for face-swapping. The comparison with previous methods on FaceForensics++ dataset demonstrates our method achieves new state-of-art results especially in preserving target attributes, keeping the visual consistency. Also our extensive analyses provide the advantages of BlendFace for subsequent face-related research.

Potential Negative Societal Impacts. Face Swapping models are at risk of abuse, *e.g.*, deceiving face verification systems and synthesizing political speeches, that are known as “deepfake”. Therefore the vision community has been working on digital face forensics, which leads to so many promising deepfake detection approaches [5, 16, 27, 28, 46, 47, 52, 62, 71, 77, 86, 87] and a wide variety of benchmarks [19, 20, 31, 38, 43, 44, 48, 60, 88, 92]. The risk can be mitigated by proactive detection methods [78, 81] and by strictly gating the release of our model only for research purpose. In addition, we will release the benchmark dataset of our model on FF++ for future studies of face forensics.

References

- [1] Deepfakes. <https://github.com/deepfakes/faceswap>. Accessed: 2023-03-02. 5
- [2] FaceSwap. <https://github.com/MarekKowalski/FaceSwap/>. Accessed: 2023-03-02. 5
- [3] Gaze estimation. <https://github.com/david-wb/gaze-estimation>. Accessed: 2023-03-02. 5
- [4] Hiface repo. <https://johann.wang/HifiFace/>. Accessed: 2023-03-02. 5
- [5] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen. MesoNet: a Compact Facial Video Forgery Detection Network. In *WIFS*, 2018. 8
- [6] Xiang An, Xuhan Zhu, Yuan Gao, Yang Xiao, Yongle Zhao, Ziyong Feng, Lan Wu, Bin Qin, Ming Zhang, Debing Zhang, et al. Partial FC: Training 10 Million Identities on a Single Machine. In *ICCV*, 2021. 2, 3
- [7] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In *ICML*, 2017. 1, 2
- [8] Sven Bambach, Stefan Lee, David J. Crandall, and Chen Yu. Lending A Hand: Detecting Hands and Recognizing Activities in Complex Egocentric Interactions. In *ICCV*, 2015. 2
- [9] Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, and Gang Hua. Towards Open-Set Identity Preserving Face Synthesis. In *CVPR*, 2018. 1
- [10] Dmitri Bitouk, Neeraj Kumar, Samreen Dhillon, Peter Belhumeur, and Shree K. Nayar. Face Swapping: Automatically Replacing Faces in Photographs. *TOG*, 27(3), 2008. 2
- [11] Volker Blanz, Kristina Scherbaum, Thomas Vetter, and Hans-Peter Seidel. Exchanging Faces in Images. *CGF*, 23, 2004. 2
- [12] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. VGGFace2: A dataset for recognising faces across pose and age. In *FG*, 2018. 2, 3, 5
- [13] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], 2015. 2
- [14] Renwang Chen, Xuanhong Chen, Bingbing Ni, and Yanhao Ge. SimSwap: An Efficient Framework For High Fidelity Face Swapping. In *ACMMM*, 2020. 1, 2, 3, 5
- [15] Yi-Ting Cheng, Virginia Tzeng, Yu Liang, Chuan-Chang Wang, Bing-Yu Chen, Yung-Yu Chuang, and Ming Ouhyoung. 3D-Model-Based Face Replacement in Video. In *SIGGRAPH*, 2009. 2
- [16] Davide Cozzolino, Andreas Rössler, Justus Thies, Matthias Nießner, and Luisa Verdoliva. ID-Reveal: Identity-aware DeepFake Video Detection. In *ICCV*, 2021. 8
- [17] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive Angular Margin Loss for Deep Face Recognition. In *CVPR*, 2019. 1, 2, 3, 4, 5, 7
- [18] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3D Face Reconstruction with Weakly-Supervised Learning: From Single Image to Image Set. In *CVPRW*, 2019. 2, 3, 5
- [19] Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. The deepfake detection challenge dataset. *arXiv:2006.07397*, 2020. 8
- [20] Brian Dolhansky, Russ Howes, Ben Pflaum, Nicole Baram, and Cristian Canton Ferrer. The deepfake detection challenge (dfdc) preview dataset. *arXiv:1910.08854*, 2019. 8
- [21] Alireza Fathi, Xiaofeng Ren, and James M. Rehg. Learning to recognize objects in egocentric activities. In *CVPR*, 2011. 2
- [22] Gege Gao, Huaibo Huang, Chaoyou Fu, Zhaoyang Li, and Ran He. Information bottleneck disentanglement for identity swapping. In *CVPR*, 2021. 1, 2, 3, 5
- [23] Sixue Gong, Xiaoming Liu, and Anil K Jain. Jointly debiasing face recognition and demographic attribute estimation. In *ECCV*, 2020. 3
- [24] Sixue Gong, Xiaoming Liu, and Anil K Jain. Mitigating face recognition bias via group adaptive classifier. In *CVPR*, 2021. 3
- [25] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 1, 2
- [26] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *ECCV*, 2016. 2, 3, 4, 5
- [27] Alexandros Haliassos, Rodrigo Mira, Stavros Petridis, and Maja Pantic. Leveraging real talking faces via self-supervision for robust forgery detection. In *CVPR*, 2022. 8
- [28] Alexandros Haliassos, Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Lips don't lie: A generalisable and robust approach to face forgery detection. In *CVPR*, 2021. 8
- [29] Z. He, W. Zuo, M. Kan, S. Shan, and X. Chen. Attgan: Facial attribute editing by only changing what you want. *TIP*, 2019. 1
- [30] Gary B. Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. In *Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition*, 2008. 7
- [31] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. DeeperForensics-1.0: A large-scale dataset for real-world face forgery detection. In *CVPR*, 2020. 8
- [32] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 3
- [33] Krishnapriya K S, Vítor Albiero, Kushal Vangara, Michael King, and Kevin Bowyer. Issues related to face recognition accuracy varying based on race and skin tone. *TTS*, 2020. 2
- [34] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *ICLR*, 2018. 1, 2

- [35] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *NeurIPS*, 2021. 2
- [36] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 1, 2, 5
- [37] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 2
- [38] Hasam Khalid, Shahroz Tariq, Minha Kim, and Simon S. Woo. FakeAVCeleb: A novel audio-video multimodal deepfake dataset. In *NeurIPS Datasets and Benchmarks Track (Round 2)*, 2021. 8
- [39] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. 2
- [40] Jiseob Kim, Jihoon Lee, and Byoung-Tak Zhang. Smoothswap: A simple enhancement for face-swapping with smoothness. In *CVPR*, 2022. 2, 5
- [41] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *CVPR*, 2022. 2
- [42] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [43] Pavel Korshunov and Sébastien Marcel. Deepfakes: a new threat to face recognition? assessment and detection. *arXiv preprint arXiv:1812.08685*, 2018. 8
- [44] Patrick Kwon, Jaeseong You, Gyuhyeon Nam, Sungwoo Park, and Gyeongsu Chae. Kodf: A large-scale korean deepfake detection dataset. In *ICCV*, 2021. 8
- [45] Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, and Fang Wen. Advancing high fidelity identity swapping for forgery detection. In *CVPR*, 2020. 1, 2, 3, 4, 5, 8
- [46] Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Baining Guo. Face x-ray for more general face forgery detection. In *CVPR*, 2020. 3, 8
- [47] Y. Li, M. Chang, and S. Lyu. In ictu oculi: Exposing ai created fake videos by detecting eye blinking. In *WIFS*, 2018. 8
- [48] Yuezun Li, Pu Sun, Honggang Qi, and Siwei Lyu. Celeb-df: A large-scale challenging dataset for deepfake forensics. In *CVPR*, 2020. 8
- [49] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017. 2
- [50] Yuchen Luo, Junwei Zhu, Keke He, Wenqing Chu, Ying Tai, Chengjie Wang, and Junchi Yan. Styleface: Towards identity-disentangled face generation on megapixels. In *ECCV*, 2022. 1, 2
- [51] Xudong Mao, Qing Li, Haoran Xie, Raymond Y.K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017. 1, 2
- [52] Iacopo Masi, Aditya Killekar, Royston Marian Mascarenhas, Shenoy Pratik Gurudatt, and Wael AbdAlmageed. Two-branch recurrent network for isolating deepfakes in videos. In *ECCV*, 2020. 8
- [53] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: The first manually collected, in-the-wild age database. In *CVPRW*, 2017. 7
- [54] Yuval Nirkin, Yosi Keller, and Tal Hassner. FSGAN: Subject agnostic face swapping and reenactment. In *ICCV*, 2019. 2, 5
- [55] Yuval Nirkin, Iacopo Masi, Anh Tran Tuan, Tal Hassner, and Gerard Medioni. On Face Segmentation, Face Swapping, and Face Perception. In *FG*, 2018. 2
- [56] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 4
- [57] Haibo Qiu, Baosheng Yu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. Synface: Face recognition with synthetic data. In *ICCV*, 2021. 7
- [58] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv:1511.06434*, 2015. 1, 2
- [59] Joseph P Robinson, Can Qin, Yann Henon, Samson Timoner, and Yun Fu. Balancing Biases and Preserving Privacy on Balanced Faces in The Wild. *arXiv preprint arXiv:2103.09118*, 2021. 2
- [60] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Niessner. Faceforensics++: Learning to detect manipulated facial images. In *ICCV*, 2019. 2, 5, 8
- [61] Soumyadip Sengupta, Jun-Cheng Chen, Carlos Castillo, Vishal M. Patel, Rama Chellappa, and David W. Jacobs. Frontal to profile face verification in the wild. In *WACV*, 2016. 7
- [62] Kaede Shiohara and Toshihiko Yamasaki. Detecting deepfakes with self-blended images. In *CVPR*, 2022. 8
- [63] Nisha Srinivas, Karl Ricanek, Dana Michalski, David S. Bolme, and Michael King. Face recognition algorithm bias: Performance differences on images of children and adults. In *CVPRW*, 2019. 2
- [64] Philipp Terhörst, Jan Niklas Kolf, Marco Huber, Florian Kirchbuchner, Naser Damer, Aythami Morales Moreno, Julian Fierrez, and Arjan Kuijper. A Comprehensive Study on Face Recognition Biases Beyond Demographics. *TTS*, 3(1), 2021. 2
- [65] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *TOG*, 2019. 5
- [66] J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *CVPR*, 2016. 5
- [67] Ruben Vera-Rodriguez, Marta Blazquez, Aythami Morales, Ester Gonzalez-Sosa, João C. Neves, and Hugo Proença. Facegenderid: Exploiting gender information in dcnn face recognition systems. In *CVPRW*, 2019. 2
- [68] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, 2018. 2

- [69] Mei Wang and Weihong Deng. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9322–9331, 2020. 3
- [70] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *CVPR*, 2019. 3
- [71] Run Wang, Felix Juefei-Xu, Lei Ma, Xiaofei Xie, Yihao Huang, Jian Wang, and Yang Liu. Fakespotter: A simple yet robust baseline for spotting ai-synthesized fake faces. In *IJCAI*, 2020. 8
- [72] Yuhan Wang, Xu Chen, Junwei Zhu, Wenqing Chu, Ying Tai, Chengjie Wang, Jilin Li, Yongjian Wu, Feiyue Huang, and Rongrong Ji. Hiface: 3d shape and semantic prior guided high fidelity face swapping. In *IJCAI*, 2021. 1, 2, 3, 4, 5
- [73] Chao Xu, Jiangning Zhang, Yue Han, Guanzhong Tian, Xianfang Zeng, Ying Tai, Yabiao Wang, Chengjie Wang, and Yong Liu. Designing one unified framework for high-fidelity face reenactment and swapping. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *ECCV*, 2022. 2
- [74] Chao Xu, Jiangning Zhang, Miao Hua, Qian He, Zili Yi, and Yong Liu. Region-aware face swapping. In *CVPR*, 2022. 1, 2, 4
- [75] Yangyang Xu, Bailin Deng, Junle Wang, Yanqing Jing, Jia Pan, and Shengfeng He. High-resolution face swapping via latent semantics disentanglement. In *CVPR*, 2022. 1, 2, 3, 5, 6
- [76] Zhiliang Xu, Hang Zhou, Zhibin Hong, Ziwei Liu, Jiaming Liu, Zhizhi Guo, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. Styleswap: Style-based generator empowers robust face swapping. In *ECCV*, 2022. 1, 2, 4
- [77] X. Yang, Y. Li, and S. Lyu. Exposing deep fakes using inconsistent head poses. In *ICASSP*, 2019. 8
- [78] Yuankun Yang, Chenyue Liang, Hongyu He, Xiaoyu Cao, and Neil Zhenqiang Gong. Faceguard: Proactive deepfake detection. *arXiv preprint arXiv:2109.05673*, 2021. 8
- [79] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014. 2, 3
- [80] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. BiSeNet: Bilateral Segmentation Network for Real-time Semantic Segmentation. In *ECCV*, 2018. 2, 3, 4
- [81] Ning Yu, Vladislav Skripniuk, Sahar Abdelnabi, and Mario Fritz. Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In *ICCV*, 2021. 8
- [82] Seyma Yucer, Samet Akçay, Noura Al-Moubayed, and Toby P Breckon. Exploring racial bias within face recognition via per-subject adversarially-enabled data augmentation. In *CVPRW*, 2020. 2
- [83] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *ICCV*, 2019. 1
- [84] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014. 8
- [85] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 3
- [86] Tianchen Zhao, Xiang Xu, Mingze Xu, Hui Ding, Yuanjun Xiong, and Wei Xia. Learning self-consistency for deepfake detection. In *ICCV*, 2021. 8
- [87] Yinglin Zheng, Jianmin Bao, Dong Chen, Ming Zeng, and Fang Wen. Exploring temporal coherence for more general video face forgery detection. In *ICCV*, 2021. 8
- [88] Tianfei Zhou, Wenguan Wang, Zhiyuan Liang, and Jianbing Shen. Face forensics in the wild. In *CVPR*, 2021. 8
- [89] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 1, 2
- [90] Yuhao Zhu, Qi Li, Jian Wang, Chengzhong Xu, and Zhenan Sun. One shot face swapping on megapixels. In *CVPR*, June 2021. 1, 2, 3, 5, 6
- [91] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Jiwen Lu, Dalong Du, et al. Webface260m: A benchmark unveiling the power of million-scale deep face recognition. In *CVPR*, 2021. 2, 3
- [92] Bojia Zi, Minghao Chang, Jingjing Chen, Xingjun Ma, and Yu-Gang Jiang. Wilddeepfake: A challenging real-world dataset for deepfake detection. In *ACMMM*, 2020. 8