

# When Age-Invariant Face Recognition Meets Face Age Synthesis: A Multi-Task Learning Framework and a New Benchmark

Zhizhong Huang<sup>✉</sup>, Graduate Student Member, IEEE, Junping Zhang<sup>✉</sup>, Senior Member, IEEE, and Hongming Shan<sup>✉</sup>, Senior Member, IEEE

**Abstract**—To minimize the impact of age variation on face recognition, age-invariant face recognition (AIFR) extracts identity-related discriminative features by minimizing the correlation between identity- and age-related features while face age synthesis (FAS) eliminates age variation by converting the faces in different age groups to the same group. However, AIFR lacks visual results for model interpretation and FAS compromises downstream recognition due to artifacts. Therefore, we propose a unified, multi-task framework to jointly handle these two tasks, termed MTLFace, which can learn the age-invariant identity-related representation for face recognition while achieving pleasing face synthesis for model interpretation. Specifically, we propose an attention-based feature decomposition to decompose the mixed face features into two uncorrelated components—identity- and age-related features—in a spatially constrained way. Unlike the conventional one-hot encoding that achieves group-level FAS, we propose a novel identity conditional module to achieve identity-level FAS, which can improve the age smoothness of synthesized faces through a weight-sharing strategy. Benefiting from the proposed multi-task framework, we then leverage those high-quality synthesized faces from FAS to further boost AIFR via a novel selective fine-tuning strategy. Furthermore, to advance both AIFR and FAS, we collect and release a large cross-age face dataset with age and gender annotations, and a new benchmark specifically designed for tracing long-missing children. Extensive experimental results on five benchmark cross-age datasets demonstrate that MTLFace yields superior performance than state-of-the-art methods for both AIFR and FAS. We further validate MTLFace on two popular general face recognition datasets, obtaining competitive performance on face recognition in the wild. The source code and datasets are available at <http://hzzzone.github.io/MTLFace>.

**Index Terms**—Face recognition, face aging, generative adversarial networks

## 1 INTRODUCTION

FACE recognition has been a hot research topic in computer vision for many years. Recently, deep-learning-based methods have achieved excellent performance, even surpassing humans in several scenarios, by empowering face recognition models with deep neural networks [1], [2], [3]. The traditional wisdom is to train the face recognition

models to increase the intra-class compactness with a massive amount of data and margin-based metrics for improved recognition performance [4], [5].

Despite the remarkable success of general face recognition (GFR), *how to minimize the impact of age variation* is a lingering challenge for current face recognition systems in correctly identifying faces in many practical applications such as tracing long-missing children. Therefore, it is of great significance to achieve face recognition without age variation, i.e., age-invariant face recognition or AIFR. AIFR, however, remains extremely challenging in the following three aspects. First, when the age gap becomes large in cross-age face recognition, age variation can dominate the facial appearance, which then significantly compromises the face recognition performance. Second, face age synthesis (FAS) is a complex process involving face aging/rejuvenation (*a.k.a* age progression/regression) since the facial appearance changes dramatically over a long time and differs from person to person. Last, it is infeasible to obtain a large-scale paired face dataset for training a model to render faces with natural effects while preserving identities.

To address the issues mentioned above, current methods for AIFR can be roughly summarized into two categories: generative and discriminative models. Given a face image, generative models [6], [7], [8] aim to transform faces of different age groups into the same age group in order to minimize the impact of age variation on face recognition. Recently, generative adversarial networks (GANs) [9] have

- Zhizhong Huang and Junping Zhang are with the Shanghai Key Lab of Intelligent Information Processing and the School of Computer Science, Fudan University, Shanghai 200433, China. E-mail: {zzhuang19, jpzhang}@fudan.edu.cn.
- Hongming Shan is with the Institute of Science and Technology for Brain-inspired Intelligence and MOE Frontiers Center for Brain Science and Key Laboratory of Computational Neuroscience and Brain-Inspired Intelligence (Ministry of Education), Fudan University, Shanghai 200433, China, and also with the Shanghai Center for Brain Science and Brain-inspired Technology, Shanghai 201210, China. E-mail: hmshan@fudan.edu.cn.

Manuscript received 3 February 2022; revised 16 October 2022; accepted 25 October 2022. Date of publication 28 October 2022; date of current version 5 May 2023.

This work was supported in part by Shanghai Municipal Science and Technology Major Project under Grant 2018SHZDZX01, in part by ZJ Lab, in part by the National Natural Science Foundation of China under Grants 62176059 and 62101136, in part by the Shanghai Municipal of Science and Technology Project under Grant 20JC1419500, and in part by Shanghai Center for Brain Science and Brain-inspired Technology.

(Corresponding author: Hongming Shan.)

Recommended for acceptance by W. Scheirer.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TPAMI.2022.3217882>, provided by the authors.

Digital Object Identifier no. 10.1109/TPAMI.2022.3217882

been successfully used to enhance the image quality of synthesized faces [10], [11], [12], [13], [14], which typically use one-hot encoding to specify the target age group. However, the one-hot encoding represents the age group-level face transformation, ignoring identity-level personalized patterns and leading to unexpected artifacts. As a result, the performance of AIFR cannot be significantly improved due to the unpleasing synthesized faces and unexpected change in identity. On the other hand, discriminative models [15], [16] focus on extracting age-invariant features by disentangling the identity-related information from mixed face information so that only the identity-related information is used by the face recognition systems. Although these models achieve promising performance in AIFR, they cannot provide users, for example policemen, with visual results just as the generative methods do to further verify identities, which can compromise model interpretability in the decision-making processes of many practical applications.

To further improve the image quality of generative models and provide model interpretability for discriminative models, in this paper, we propose a unified, multi-task learning framework to simultaneously achieve AIFR and FAS, termed MTLFace, which can enjoy the best of both worlds; i.e., learning age-invariant identity-related representations while achieving pleasing face synthesis. More specifically, we propose an attention-based feature decomposition to decompose the mixed high-level features into two uncorrelated components—identity- and age-related features—in a spatially constrained way. We then decorrelate these two components in a multi-task learning framework, in which an age estimation task is to extract age-related features while a face recognition task is to extract identity-related features; in addition, a continuous cross-age discriminator with a gradient reversal layer [17] further encourages extracting identity-related age-invariant features. Moreover, we propose an identity conditional module to achieve identity-level transformation patterns for FAS, with a weight-sharing strategy to improve the age smoothness of the synthesized faces; i.e., the faces are aged smoothly. Extensive experimental results demonstrate the superior performance of the proposed MTLFace over existing state-of-the-art methods for AIFR and FAS, and competitive performance for general face recognition in the wild. Fig. 1 presents an example of age progression and regression of the same person by our MTLFace, showing that our framework can synthesize photorealistic faces while preserving identity.

The contributions of this paper are summarized as follows.

- 1) We propose a unified, multi-task learning framework to jointly handle AIFR and FAS, which can learn age-invariant identity-related representations while achieving pleasing face synthesis.
- 2) We propose an attention-based feature decomposition to separate the age- and identity-related features on high-level feature maps, which can spatially constrain the decomposition process in contrast to the previous unconstrained decomposition on feature vectors. Age estimation and face recognition tasks are incorporated to supervise the decomposition

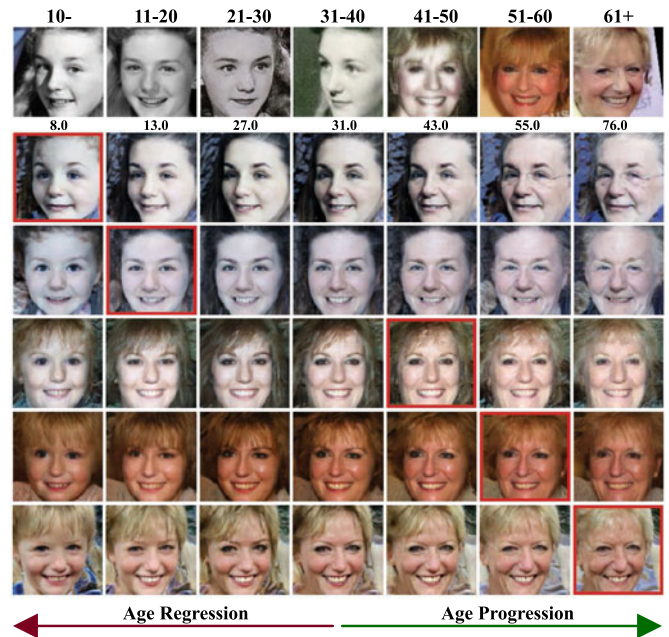


Fig. 1. Sample results of our MTLFace. First row: real faces of the same person at different ages with estimated age labels underneath. Remaining rows: synthesized faces where the input faces are given in the red boxes.

process in conjunction with a continuous domain adaptation.

- 3) We propose a novel identity conditional module to achieve identity-level face transformation, which leverages a weight-sharing strategy to improve the age smoothness of the synthesized faces in contrast to previous one-hot encoding that can only achieve age group-level face transformation. We further extend it into a multi-level architecture for improved face age synthesis.
- 4) We propose a selective fine-tuning strategy to further boost AIFR by automatically selecting those high-quality synthesized faces from FAS for fine-tuning.
- 5) We collect and release a large-scale cross-age dataset of millions of faces with balanced age and gender annotations, which can not only advance the development of the AIFR and FAS but also be useful for other face-related research tasks; e.g., pretraining for face age estimation.
- 6) To promote the utility in tracing long-missing children, we construct a new benchmark with paired child and adult faces of the same person, which contains the same identities as the Labeled Faces in the Wild (LFW) and is thus orthogonal to the current literature for future evaluation of cross-age face recognition.
- 7) Extensive experimental results demonstrate the effectiveness of the proposed framework for AIFR and FAS on five benchmark datasets, and competitive performance on two popular GFR datasets.

We note that a preliminary version of this work was published in IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2021 [18]. In this paper, we further extend our preliminary work with the following major

improvements. First, we propose a selective fine-tuning strategy that selects the high-quality synthesized faces from FAS through a two-component Gaussian mixture model and then fine-tunes the model to further boost AIFR, bringing consistent performance improvements on all benchmark datasets. Second, we extend the model by incorporating multi-level age conditional information with a StyleGAN-based architecture, which, as a result, can synthesize more photorealistic faces than previous model [18] and achieve continuous face age synthesis. Third, we construct a new benchmark testing set for the AIFR task with an emphasis on tracing long-missing children, which contains a subset of identities of LFW with manually annotated and paired child-adult faces; extensive baselines including humans and other deep-learning based face recognition models were conducted to evaluate the new benchmark. Fourth, we further conduct more experiments with more competitive baselines to evaluate the face age synthesis task of our model including the results on the new benchmark. Finally, we evaluate the model through both quantitative and qualitative comparisons between aged/rejuvenated faces and ground-truth faces to further validate the effectiveness of the proposed method. We also visualize the training process of MTLFace to show its training stability. Note that a detailed comparison between the current and preliminary versions of MTLFace can be found in supplementary [Section 1](#); our new model consistently improves the performance of face recognition and the quality of synthesized faces.

The remainder of this paper is organized as follows. A brief review of the related work on age-invariant face recognition and face age synthesis is given in [Section 2](#). We present the attention-based feature decomposition, identity conditional module, and selective fine-tuning strategy as well as the proposed MTLFace in [Section 3](#). Comprehensive experimental results on the age-invariant face recognition and face age synthesis are reported and analyzed in [Section 4](#). Finally, this paper is discussed in [Section 5](#), followed by a concluding summary in [Section 6](#).

## 2 RELATED WORK

This section briefly surveys developments in age-invariant face recognition and face age synthesis.

### 2.1 Age-Invariant Face Recognition (AIFR)

Prior studies usually minimize the impacts of age variation by disentangling age-invariant features from mixed features. For example, Gong et al. adopted hidden factor analysis (HFA) to factorize mixed features and then minimize the age variation in identity-related features [19]. Later, Wen et al. extended HFA to a deep learning framework with the latent factor guided convolutional neural network (LF-CNN) [20]. Zheng et al. introduced an age estimation task to guide AIFR [21]. Most recently, CNNs-based discriminative methods have achieved promising results for AIFR. To reduce the intra-class discrepancy caused by aging, Wang et al. proposed an orthogonal embedding CNN for AIFR [16], termed OE-CNN, which decomposes the facial embeddings into two orthogonal components, where the identity- and age-related features are represented as the angular and

radial directions, respectively. Similarly, Wang et al. proposed the decorrelated adversarial learning (DAL) to achieve feature decomposition in an adversarial manner under the assumption that the two components are uncorrelated [22]. Lee et al. proposed an inter-prototype loss to minimize the similarity between child faces [23]. Hou et al. [24] and Xie et al. [25] proposed to minimize the mutual information between the identity- and age-related components of the face image from the same person to reduce the effect of age variations.

The work related to ours is [26], in which a cGANs-based model with cross-age domain adversarial training to extract age-invariant representations, is adopted to perform the two tasks simultaneously. However, it generates over-smoothed faces with subtle changes. Different from [26], our framework has the following advantages: 1) our feature decomposition is performed on feature maps through an attention mechanism; 2) a continuous domain adaptation with a gradient reversal layer is used to learn age-invariant identity-related representation; 3) the proposed identity conditional module can achieve identity-level face synthesis and improve the age smoothness of synthesized faces; and 4) the proposed FT-Sel can further boost the performance of AIFR with the automatically selected high-quality synthesized faces from FAS.

### 2.2 Face Age Synthesis (FAS)

Existing methods for FAS can be roughly divided into physical model-, prototype-, and deep generative model-based methods.

Physical model-based methods [27], [28], [29] mechanically model the changes in faces over time, but they are computationally expensive and require massive paired images of the same person over a long period of time. Prototype-based methods [30], [31] achieve face aging/rejuvenation using the averaged faces in each age group, and consequently, the identity cannot be well preserved.

In contrast to the above two categories, deep generative model-based methods [32], [33] exploit a deep neural network for this task. For example, Wang et al. proposed recurrent face aging (RFA), which uses a recurrent neural network to model the intermediate transition states of age progression/regression, traversing on which a smooth face aging process can be achieved [33]. Inspired by the powerful capability of generative adversarial networks (GANs) [9], especially conditional GANs (cGANs) [34], in generating high-quality images, many recent studies [12], [13], [14], [35], [36], [37], [38], [39] resort to them to improve the visual quality of synthesized faces and train models with unpaired face images. For example, Zhang et al. used a conditional adversarial autoencoder (CAAE) to achieve both age progression/regression by traversing on a low-dimensional face manifold [14]. Wang et al. introduced the perceptual loss to preserve identities during face aging/rejuvenation [12]. Yang et al. designed a discriminator with a pyramid architecture to enhance aging details [13], [40].

However, these methods mainly aim to improve the visual quality of generated faces, and hardly improve the performance of AIFR due to the artifacts resulting from group-level face transformation, and the unexpected

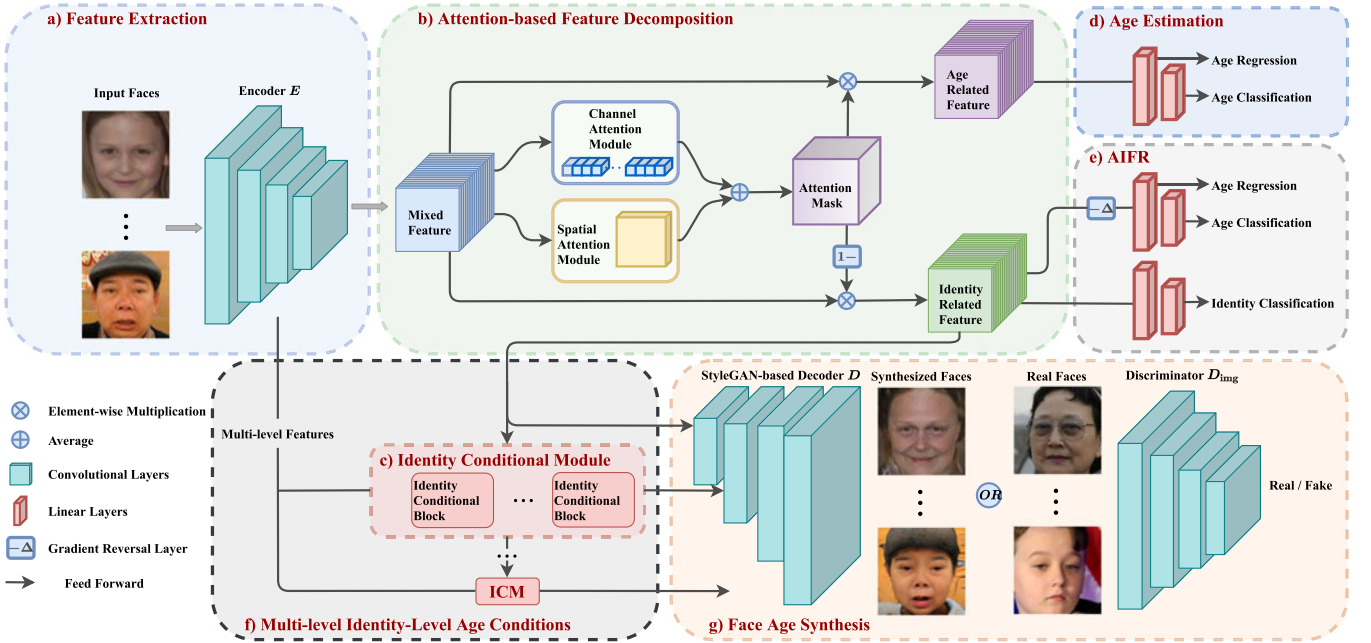


Fig. 2. An overview of the proposed MTLFace including two tasks. AIFR: The encoder  $E$  first extracts the mixed feature maps from input faces, which are then decomposed into two disjoint identity- and age-related feature maps by the multi-task training and continuous domain adaptation. FAS: The decoder  $D$  produces synthesized faces from identity-level features under the control of multi-level age conditions from the output of the identity conditional modules; the discriminator  $D_{img}$  penalizes the framework to obtain better visual quality.

change in identity. Our method differs in the following aspects: 1) the proposed MTLFace achieves AIFR and FAS simultaneously to enhance the visual quality with identity-related information from AIFR; 2) the proposed identity conditional module (ICM) achieves an identity-level face age synthesis in contrast to the previous group-level face age synthesis; 3) a weight-sharing strategy in ICM can improve the age smoothness of the synthesized faces; and 4) a well-designed generator with feature-pyramid and StyleGAN architectures is employed for photorealistic and continuous synthesis.

### 3 METHODOLOGY

Fig. 2 presents the architecture of the proposed MTLFace, which is detailed in the following subsections.

#### 3.1 Attention-Based Feature Decomposition

As the faces age over time, the critical problem of AIFR is that the age variation usually increases intra-class distances. As a result, it becomes challenging to correctly recognize two face images of the same person with a large age gap, since the mixed facial representations are severely entangled with unrelated information such as facial shape and texture changes. Recently, Wang et al. designed a linear factorization module to decompose feature vectors into two unrelated components [22]. Formally, the feature vector  $\mathbf{x} \in \mathbb{R}^d$ , extracted from an input image  $I \in \mathbb{R}^{3 \times H \times W}$ , can be decomposed as [22]:

$$\mathbf{x} = \mathbf{x}_{age} + \mathbf{x}_{id}, \quad (1)$$

where  $\mathbf{x}_{age}$  and  $\mathbf{x}_{id}$  denote the age- and identity-related components, respectively. This decomposition is implemented through a residual mapping. However, it has the following

drawbacks: 1) this decomposition performs on one-dimensional feature vector, and the resultant identity-related component lacks spatial information of face, not suitable for FAS; and 2) this decomposition is unconstrained, which may lead to unstable training.

To address these drawbacks, we instead propose to decompose the mixed feature-maps in a high-level semantic space through an attention mechanism, termed attention-based feature decomposition or AFD. The main reason is that manipulating on the feature vectors is more complicated than on the feature maps since the aging/rejuvenation effects, such as beards and wrinkles, appear in the semantic feature space but are lost in one-dimensional feature vectors. Formally, we use a ResNet-like backbone as encoder  $E$  to extract mixed feature maps  $\mathbf{X} \in \mathbb{R}^{C \times H' \times W'}$  from an input image  $I$ , i.e.  $\mathbf{X} = E(I)$ , the AFD can be expressed as:

$$\mathbf{X} = \underbrace{\mathbf{X} \circ \sigma(\mathbf{X})}_{\mathbf{x}_{age}} + \underbrace{\mathbf{X} \circ (1 - \sigma(\mathbf{X}))}_{\mathbf{x}_{id}}, \quad (2)$$

where  $\circ$  denotes element-wise multiplication and  $\sigma$  represents an attention module. In doing so, the age-related information in the feature maps can be separated through the attention module supervised by an age estimation task, and the residual part, regarded as the identity-related information, can be supervised by a face recognition task. Consequently, the attention mechanism constrains the decomposition module, better at detecting the age-related features in semantic feature maps. We note that  $\mathbf{X}$  is assumed to only contain the age and identity information as driven by the two corresponding tasks, and the remaining information such as background is important for FAS, which is preserved by skip connections from encoder to decoder. Fig. 2b details the proposed AFD.




Fig. 3. Comparison between one-hot encoding and ICB.

In this paper, we adopt the average of channel attention (CA) [41] and spatial attention (SA) [42] to highlight age-related information at both channel and spatial levels. Note that the outputs of these two attentions have different sizes; we first stretch each of them to the original input size and then average them. Different attention modules such as CA, SA, and CBAM [42], are also investigated in Section 4.

### 3.2 Identity Conditional Module

The mainstream face aging studies [10], [11], [12], [13], [14] usually split ages into several non-overlapping age groups, since the changes over time are minor with a small age gap. These methods typically use one-hot encoding to specify the target age group to control the aging/rejuvenation process [10], [12], [14] as illustrated in Fig. 3a. Consequently, a group-level aging/rejuvenation pattern, such as having a beard at age 30, is learned for each age group due to the use of one-hot age conditions. This has two drawbacks: 1) the one-hot encoding represents the age group-level aging/rejuvenation pattern, ignoring identity-level personalized patterns, particularly for different genders and races; and 2) one-hot encoding may not ensure the age smoothness of synthesized faces.

To address these issues raised by one-hot encoding, we propose an identity conditional block (ICB) to achieve identity-level aging/rejuvenation pattern, with a weight-sharing strategy to improve the age smoothness of synthesized faces. Specifically, the proposed ICB takes the identity-related feature from AFD as input to learn an identity-level aging/rejuvenation pattern. Note that the input features for ICB may be mixed up with the ages due to the use of unpaired training data, which could harm the learning of identity-level pattern. Instead, in our multi-task learning framework, under the supervision from face classification and domain adaptation loss, the identity features are invariant to ages. As a result, the ages would not have an impact on the training of ICB. Next, we propose a weights-sharing strategy to improve the age smoothness of synthesized faces so that some convolutional filters are shared across adjacent age groups as shown in Fig. 3b. The rationale behind this idea is that faces gradually change over time, and the shared filters can learn some common aging/rejuvenation patterns between adjacent age groups. Note that the number of channels of the input feature maps is reduced to 1/4 of the

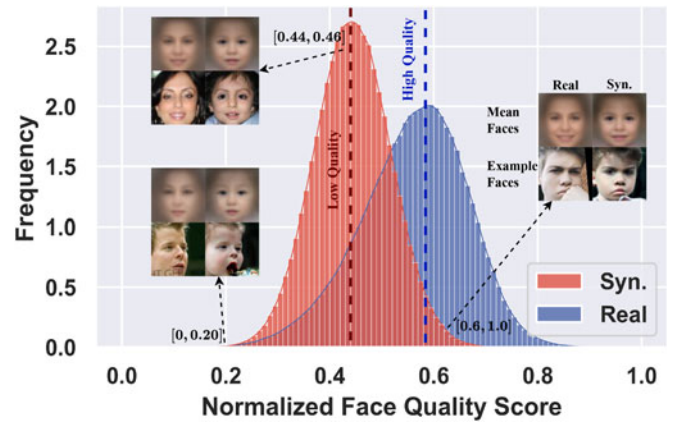


Fig. 4. Visualization of the normalized face quality scores of synthesized and real faces, where the two vertical dotted lines denote the mean values of two components of the GMM, and a higher value indicates better quality. We showcase some examples with their rejuvenated faces, and the mean faces at corresponding scores.

original ones using  $1 \times 1$  convolutions to reduce the computational cost. In this paper, a hyper-parameter  $s$  is used control how many filters are shared for two adjacent age groups, which is empirically set to  $1/8$ ; i.e., two adjacent age groups share 16 filters. We stack ICBs to form an identity conditional module (ICM), as shown in Fig. 2c.

### 3.3 Selective Fine-Tuning Strategy

In this section, we present a novel selective fine-tuning strategy (FT-Sel), which can *selectively* include high-quality synthesized faces into the training data to further boost the performance of face recognition for AIFR task.

Our motivation comes from the observation that there are a few child faces below 10 years old in the training data since it is cumbersome to collect the paired child and adult faces. Such imbalanced data (i.e. lacking child faces) would significantly harm the face recognition performance on child's faces, affecting the practical application for tracking long-missing children. As a result, the intra-class distances between children are not discriminative, which has been widely reported in long-tailed learning [43].

Benefiting from the proposed multi-task framework, we address this problem by transforming the faces above 10 years old into 10- to obtain the paired child and adult faces. This is similar to [44] that leverages style transfer [45] to augment the training data. If MTLFace disentangles the faces to identity (i.e. content) and age (i.e. style), we are able to alter the age styles at low-level visual images while preserving the semantic identities. As a result, the face recognition model can be encouraged to learn style-invariant identity features to further boost performance. However, only high-quality synthesized faces should be considered, as the ones with unexpected artifacts would harm the performance of face recognition [46]. To this end, we propose a selective fine-tuning strategy (FT-Sel), which can automatically select the high-quality synthesized faces using the face quality scores. Fig. 4 visualizes the face quality scores [5] of synthesized and real faces, where a significant discrepancy between the two kinds of faces can be observed. The synthesized faces with the low face quality scores (e.g. [0, 0.2]) present strong artifacts. When the scores increase, the synthesized faces become

more photorealistic despite some occlusion (e.g. [0.6,1.0]). The corresponding mean face reveals more details at the same time. This implies that the face quality scores can measure the quality of synthesized faces. As a result, the synthesized child faces of better face quality can be selected to construct more balanced training data for further boosting the performance of face recognition.

A natural question arises: *how to select those high-quality faces?* A straightforward way is to set a threshold, where the faces above this threshold can be considered as high-quality samples. However, the introduced thresholding hyperparameter needs to be *manually* tuned for satisfied recognition performance. By visualizing the distributions of the scores in Fig. 4, we propose to model the distribution of face quality scores of synthesized and real faces by a two-component Gaussian Mixture Model (GMM) for *automatically* selecting high-quality faces, which is defined as follows:

$$p(s) = \sum_{c=0}^1 p(s, c) = \sum_{c=0}^1 p(c) p(s|c), \quad (3)$$

where  $s$  denotes the face quality score, and  $c$  is a latent variable: the cluster  $c = 1$  with a higher mean face-quality score denotes the faces with higher quality while  $c = 0$  corresponds to the faces with lower quality, as illustrated in Fig. 4. Through the two-component GMM, we can infer the posterior probability of one synthesized face being of high quality, and obtain the photorealistic faces by  $p(c = 1|s) > p(c = 0|s)$  from all synthesized faces.

After selecting the high-quality synthesized faces, we then fine-tune the last linear layer for AIFR while keeping the convolutional layers fixed. In long-tailed learning, rebalancing data distribution may lead to unstable training. We highlight that FT-Sel would not change the data distribution too much as there are only 10% synthesized faces in new training data. If this problem indeed happens, one can turn to the distribution alignment technique [47] in long-tailed learning to stabilize the training. We also note that this step would not influence FAS since the last linear layer was not involved in the synthesizing process of our multi-task learning framework, which will be detailed in next section.

### 3.4 Multi-Task Learning Framework

This section presents our MTLFace including AIFR and FAS.

#### 3.4.1 Age-Invariant Face Recognition (AIFR) Task

To encourage AFD to robustly decompose features, we use an age estimation task and a face recognition task to supervise the feature decomposition. Specifically,  $\mathbf{X}_{\text{age}}$  draws the age variations through an age estimation task while  $\mathbf{X}_{\text{id}}$  encodes the identity-related information. First, as illustrated in Fig. 2d we include an age estimation network  $A$  with two linear layers of 512 and 101 neurons to achieve age regression similar to deep expectation (DEX) [48], [49], which learns the age distribution by computing a softmax expected value. Second, we append another linear layer  $\mathbf{W} \in \mathbb{R}^{101 \times n_g}$  on top of  $A$  for age classification, regularizing the learned distribution, where  $n_g$  denotes the number of age groups. The loss function to optimize age estimation

can be defined as:

$$\begin{aligned} \ell_{\text{AE}}(\mathbf{X}_{\text{age}}) = & \mathbb{E}_{\mathbf{I}}[\ell_{\text{MSE}}(\text{DEX}(A(\mathbf{X}_{\text{age}})), y_{\text{age}}) \\ & + \ell_{\text{CE}}(A(\mathbf{X}_{\text{age}})\mathbf{W}, c_{\text{age}})], \end{aligned} \quad (4)$$

where  $y_{\text{age}}$ ,  $c_{\text{age}}$ ,  $\ell_{\text{MSE}}$ , and  $\ell_{\text{CE}}$  are the ground truth age, ground truth age group, mean squared error (MSE) for age regression, and cross-entropy (CE) loss for age group classification, respectively.

Next, we leverage one linear layer  $L$  of 512 neurons to extract the feature vectors, and use the CosFace loss to supervise the learning of  $\mathbf{X}_{\text{id}}$  for identity classification. We also introduce a cross-age domain adversarial learning that encourages  $\mathbf{X}_{\text{id}}$  to be age-invariant through a continuous domain adaptation [50] with a gradient reversal layer (GRL) [17]. Fig. 2e shows these two components. The final loss for AIFR is formulated as:

$$\begin{aligned} \mathcal{L}^{\text{AIFR}} = & \ell_{\text{COSFACE}}(L(\mathbf{X}_{\text{id}}), y_{\text{id}}) \\ & + \lambda_{\text{age}}^{\text{AIFR}} \mathcal{L}_{\text{AE}}(\mathbf{X}_{\text{age}}) + \lambda_{\text{id}}^{\text{AIFR}} \mathcal{L}_{\text{AE}}(\text{GRL}(\mathbf{X}_{\text{id}})), \end{aligned} \quad (5)$$

where the first term is the CosFace loss; the second term is the age estimation loss; the last term is the domain adaptation loss;  $y_{\text{id}}$  is the identity label; and  $\lambda_*$  controls the balance of different loss terms. Note that the second and third terms use the same network structure but have different inputs and are trained independently. The activation functions and batch normalizations are ignored for simplicity, and our face recognition model is designed strictly following the setting in [51] except the AFD.

#### 3.4.2 Face Age Synthesis (FAS) Task

Figs. 2f and 2g demonstrate the FAS process of our proposed method. In detail, the decoder  $D$  adopts the architecture of the StyleGAN-based generator [52], [53], which receives the discriminative facial representations  $\mathbf{X}_{\text{id}}$  at a coarse level. Then, the single-level identity-level age condition is derived from the discriminative facial representations  $\mathbf{X}_{\text{id}}$  and the same level high-resolution features extracted from the encoder  $E$  using two ResBlocks [1]. Similar to the feature pyramid network [54], the multi-level identity-level age conditions are formed by stacking multiple ICMs at different levels. Finally, the decoder  $D$  reconstructs the progressed/regressed faces from  $\mathbf{X}_{\text{id}}$  under the control of the multi-level learned identity-level age conditions, using two StyleGAN block [52], [53] at each level with AdaIN normalization layers [55]. Formally, the process of rendering the input face  $\mathbf{I}$  to the synthesized face  $\hat{\mathbf{I}}_t$  that belongs to the target age group  $t$  can be written as:

$$\mathbf{C}_t^1 = f_1(\mathbf{X}_{\text{id}}, \mathbf{E}^1, t), \quad (6)$$

$$\mathbf{C}_t^l = f_l(\mathbf{C}_t^{l-1}, \mathbf{E}^l, t), \quad l \in \{2, 3\} \quad (7)$$

$$\hat{\mathbf{I}}_t = D(\mathbf{X}_{\text{id}}, \{\mathbf{C}_t^l\}_{l=1}^3), \quad (8)$$

where  $l$  denotes the index of different levels, and  $\mathbf{C}_l$  and  $f_l$  are the identity-level age condition and ICMs at the  $l$ th level, respectively. We note that  $D$  employs several convolutional layers at each level to map the age conditions into the style latent space of StyleGAN. The rationale is that the coarse-

level of  $\mathbf{X}_{\text{id}}$  corresponds to the shape of faces, i.e. identities, while the fine-level of age conditions renders the faces with detail aging/rejuvenation effects.

To facilitate the visual quality of the generated faces, the FAS task is trained using GANs framework. In this paper, we adopt a combination of PatchDiscriminator [56] and a StyleGAN-based discriminator [52], [53] as our discriminator  $D_{\text{img}}$  to emphasize the local-patches of generated and real images. Furthermore, the least-squares GANs [57] are employed to optimize the GANs framework for improved quality of the generated images and stable training process, which can be formulated as follows:

$$\mathcal{L}_{\text{adv}}^{\text{FAS}} = \frac{1}{2} \mathbb{E}_{\mathbf{I}} \left[ D_{\text{img}}(\hat{\mathbf{I}}_t, c^t) - 1 \right]^2, \quad (9)$$

where  $c^t$  is a scalar used in the traditional cGANs framework to align with the age condition by specifying the logits of  $t$ th target age group. To preserve the identities of input faces and improve the age accuracy, we leverage the encoder  $E$ , AFD, and the linear layer  $L$  to supervise the FAS task. Consequently, we can achieve both face aging and rejuvenation in a holistic, end-to-end manner, as illustrated in Fig. 2. This process can be formulated as follows:

$$\mathbf{X}_{\text{age}}^t, \mathbf{X}_{\text{id}}^t = \text{AFD}(E(\hat{\mathbf{I}}_t)), \quad (10)$$

$$\mathcal{L}_{\text{age}}^{\text{FAS}} = \ell_{\text{CE}}(A(\mathbf{X}_{\text{age}}^t) \mathbf{W}, c_{\text{age}}^t), \quad (11)$$

$$\mathcal{L}_{\text{id}}^{\text{FAS}} = \mathbb{E}_{\mathbf{X}_s} \left\| \mathbf{X}_{\text{id}}^t - \mathbf{X}_{\text{id}} \right\|_F^2 - \cos(L(\mathbf{X}_{\text{id}}^t), L(\mathbf{X}_{\text{id}})), \quad (12)$$

where  $\|\cdot\|_F$  represents the Frobenius norm, and  $\cos(\cdot)$  denotes cosine similarity. Another learned perceptual image patch similarity (LPIPS) loss [58] is employed to further maintain the perceptual consistency between input and synthesized faces:

$$\mathcal{L}_{\text{lpiPs}}^{\text{FAS}} = \mathbb{E}_{\mathbf{I}} \left[ \text{LPIPS}(\mathbf{I}_t, \hat{\mathbf{I}}_t) \right]. \quad (13)$$

The final loss to optimize this task can be written as:

$$\mathcal{L}^{\text{FAS}} = \lambda_{\text{adv}}^{\text{FAS}} \mathcal{L}_{\text{adv}}^{\text{FAS}} + \lambda_{\text{id}}^{\text{FAS}} \mathcal{L}_{\text{id}}^{\text{FAS}} + \lambda_{\text{age}}^{\text{FAS}} \mathcal{L}_{\text{age}}^{\text{FAS}} + \lambda_{\text{lpiPs}}^{\text{FAS}} \mathcal{L}_{\text{lpiPs}}^{\text{FAS}}, \quad (14)$$

where  $\lambda_{*}^{\text{FAS}}$  controls the importance of different loss terms of FAS task. The loss function to optimize the discriminator  $D_{\text{img}}$  in the context of least-squares GANs is defined as:

$$\mathcal{L}_{D_{\text{img}}}^{\text{FAS}} = \frac{1}{2} \mathbb{E}_{\mathbf{I}_t} \left[ D_{\text{img}}(\mathbf{I}_t, c^t) - 1 \right]^2 + \frac{1}{2} \mathbb{E}_{\mathbf{I}} \left[ D_{\text{img}}(\hat{\mathbf{I}}_t, c^t) \right]^2. \quad (15)$$

At the testing stage, the only difference from existing FAS methods is that our method needs to specify the corresponding group of filters. Consequently, our method enjoys the advantages similar to [59] that the computational cost can be significantly reduced by encoding input faces only once, instead of the  $n_g$  times needed in previous works [10], [11], [12], [13], [14], where  $n_g$  is the number of age groups. In addition, MTLFace can also achieve continuous face age synthesis similar to [40] thanks to the StyleGAN-based decoder by interpolating in the style latent space.

### 3.4.3 Optimization and Inference

The training algorithm of the proposed MTLFace is shown in Algorithm 1, which contains the following two stages.

In the first stage, the AIFR learns the discriminative facial representations and age estimation while the FAS produces the visual results that can boost the model interpretability for AIFR. Therefore, both two tasks can be jointly accomplished through a GAN-like optimization; they mutually leverage each other to boost themselves. In other words, the AIFR encourages FAS to render faces to preserve identity while FAS can facilitate the extraction of the identity-related features and boost the model interpretability for AIFR. Consequently, we alternately train these two tasks in a unified, multi-task, end-to-end framework. In the second stage, the well-trained model of previous stage is employed to synthesize paired faces, and those high-quality ones are selected for further fine-tuning with the proposed FT-Sel as illustrated in Section 3.3.

---

#### Algorithm 1. Training Algorithm of MTLFace

---

**Input:** Dataset  $\mathcal{D} = \{(\mathbf{I}, \mathbf{y}_{\text{age}}, \mathbf{y}_{\text{id}})\}$ ; trainable functions  $E, A, W, L, \{f_l\}_{l=1}^3, D$  and  $D_{\text{img}}$

**Output:** The full trained model of MTLFace.

**repeat**

    Training  $E, A, W$  and  $L$  with  $\mathcal{L}^{\text{AIFR}}$

    Training  $D_{\text{img}}$  with  $\mathcal{L}_{D_{\text{img}}}^{\text{FAS}}$

    Training  $\{f_l\}_{l=1}^3$  and  $D$  with  $\mathcal{L}^{\text{FAS}}$

**until** reaching max iteration

Fine-tuning  $L$  with FT-Sel

---

## 4 EXPERIMENTS

### 4.1 Data Collection

#### 4.1.1 Large-Scale Cross-Age Face Dataset

Current research on AIFR lacks a large-scale face dataset of millions of face images with large age gaps. To advance the development of AIFR and FAS, we collect and release a new large cross-age face dataset (LCAF) with 1.7M faces from cross-age celebrities. The collection process for our dataset is summarized in the following three steps. First, we use the public Azure Facial API [60] to estimate the ages and genders of faces from the clean MS-Celeb-1M dataset provided by [51]. Second, we randomly select faces from a total of 5M faces to check whether the faces are correctly labeled, and manually correct them as best we can if any apparent mistakes occur; we mainly focus on young ages, under 20 that are often mislabeled by the API [60]. Finally, a large-scale balanced age dataset is constructed by balancing both age and gender.

We further build a subset of the cross-age face dataset (SCAF) containing approximately 0.5M images from 12k individuals following [16], [22] for fair comparisons. We note that the training dataset, i.e. LCAF, and the testing datasets summarized in Table 1 may have very few, or even no identities overlapping as [51] already removed 500+ identities from their clean MS-Celeb-1M dataset by checking the similarity of faces between the training and testing data. Figs. 5a and 5b present example images and dataset statistics of SCAF.

TABLE 1  
Statistics of Testing Sets Used in This Paper for AIFR Task

Dataset	Subjects	Images	Pairs	Avg Age Gap of Test Set (years)
LFW [61]	5,749	13,233	6k	11.9
CACD [64]	2,000	163,446	4k	11.7
CALFW [62]	5,749	12,174	6k	17.6
AgeDB [63]	568	16,488	6k	16.8
FG-NET [65]	82	1,002	only for identification	
ECAF	613	5,265	6k	41.3

Due to the limited number of faces, FG-NET was mainly used for face identification. Our new benchmark dataset ECAF has larger age gaps than the other datasets from children to adults.

Following the mainstream literature [10], [11], [13], [37], [59] with the time span of 10 years for each age group, the ages in this paper are divided into seven non-overlapping groups; i.e., 10-, 11-20, 21-30, 31-40, 41-50, 51-60, and 61+. Note that it is a much more challenging problem to perform FAS on seven groups than on the four groups widely used in previous work.

#### 4.1.2 A New Benchmark

Compared to Labeled Faces in the Wild (LFW) [61] dataset used in GFR, existing benchmark cross-age face datasets for AIFR such as CALFW [62] and AgeDB [63], are designed to challenge current face recognition models with cases of large age gaps. However, these datasets overlook the fact that the most challenging case of AIFR is to match faces from children to adults, which is critical for practical applications such as tracing long-missing children.

Therefore, we build a new benchmark, termed Evaluation of Cross-Age Face Recognition (ECAF) for AIFR, which particularly focuses on child and adult faces of the same person. Specifically, by querying the names in LFW with the keyword “child”, we manually collect, annotate, and check the faces with the source link to ensure that the child and adult faces are from the same person. Then, we create two tasks based on ECAF to evaluate the performance of face recognition. The

first task involves 6k image pairs for  $\langle \text{Adult, Child} \rangle$ , where  $\langle \text{Adult, Child} \rangle$  denotes face verification between adult and child faces. These 6k pairs are split into 10 folds, and each fold consists of 300 intra-class and 300 inter-class pairs. The second task involves 4k pairs of  $\langle \text{Child, Child} \rangle$ , where the  $\langle \text{Child, Child} \rangle$  pairs are obtained with a similar procedure. Both tasks follow the protocols of LFW. We find that the second task is also challenging and worthy of further study by the face recognition community.

Table 1 summarizes the testing sets used in this paper for AIFR and Fig. 5c showcases some examples in ECAF. We obtain the estimated ages in the same way as stated in Section 4.1.1. ECAF has much larger age gaps than LFW [61], CACD-VS [64], CALFW [62], and AgeDB [63], and has more identities and faces than FG-NET [65]. In summary, there are three important properties of ECAF: 1) clean labels, as we manually check all faces and links in the dataset; 2) larger age gaps, as ECAF focuses on face recognition between children and adults; and 3) orthogonal to existing research on face recognition, as ECAF contains the same identities as LFW, which makes the results on ECAF reliable since there are no overlapping identities in current widely-used training datasets such as MS-Celeb-1M.

#### 4.2 Implementation Details

Following [51], we adopt ResNet-50 as the encoder  $E$ . In the decoder  $D$ , the identity age condition is bilinearly upsampled and processed with multi-level high-resolution features extracted from  $E$  by two ResBlocks [1], each of which is followed by instance normalization [66] and leaky ReLU activation of 0.2 negative slope. There are four ICBs in ICM. The age conditions are derived using several stride-2 convolutional layers and one linear layer. The style dimension of the StyleGAN-based decoder is 512. For the discriminator  $D_{\text{img}}$ , we build upon the StyleGAN architecture with 4 StyleBlocks, followed by spectral normalization [67] and leaky ReLU activation except in the last block, yielding an  $8 \times 8$  confidence map. AIFR is optimized by stochastic gradient descent (SGD) with an initial learning rate of 0.1 and a momentum of 0.9, while the ICM, the decoder  $D$ , and  $D_{\text{img}}$  are trained by Adam [68]

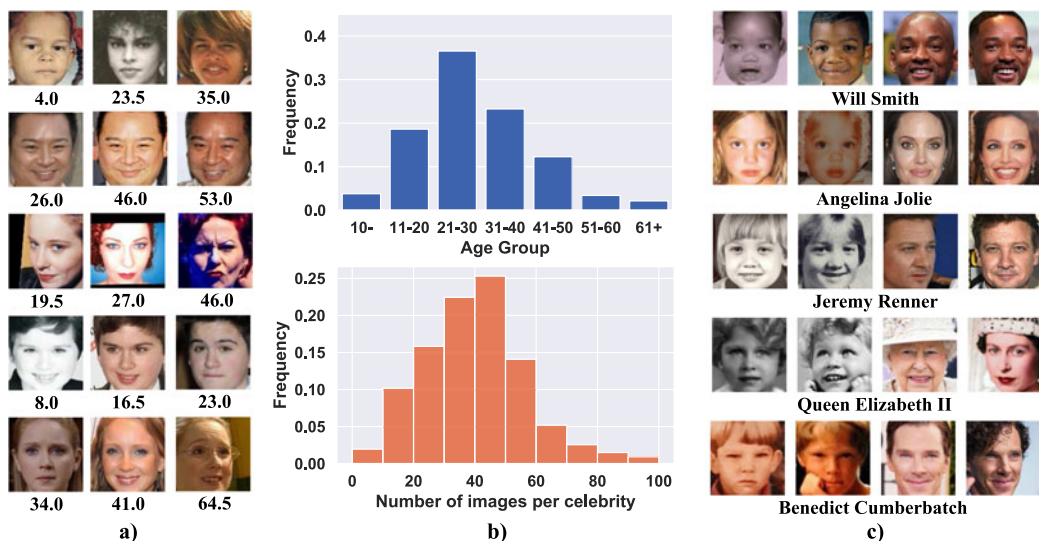


Fig. 5. a) Sample faces and b) dataset statistics on SCAF; and c) Sample faces on ECAF with celebrity names underneath.



TABLE 2  
Verification Rate (%) on AgeDB-30 Dataset for AIFR

Method	Accuracy (%)
RJIVE [71]	55.20
VGG Face [72]	89.89
Center Loss [4]	93.72
SphereFace [73]	91.70
CosFace [69]	94.56
ArcFace [51]	95.15
DAAE [37]	95.30
MTLFace (ours)	<b>96.45</b>

with a fixed learning rate of  $1.0 \times 10^{-4}$ ,  $\beta_1$  of 0.9 and  $\beta_2$  of 0.99 for face age synthesis. We train all models with a mini-batch size of 512 on 8 NVIDIA GTX 2080Ti GPUs, with 110k iterations for LCAF and 36k iterations for SCAF. The learning rate of AIFR is warmed up linearly from 0 to 0.1, and then reduced by a factor of 0.1, at iterations 5k, 70k, and 90k on LCAF and 1k, 20k, 23k on SCAF. The hyper-parameters in the loss functions are empirically set as follows:  $\lambda_{\text{age}}^{\text{AIFR}}$  is 0.001,  $\lambda_{\text{id}}^{\text{AIFR}}$  is 0.002,  $\lambda_{\text{adv}}^{\text{FAS}}$  is 5.0,  $\mathcal{L}_{\text{lipips}}^{\text{FAS}}$  is 1.0,  $\lambda_{\text{id}}^{\text{FAS}}$  is 1.0, and  $\lambda_{\text{age}}^{\text{FAS}}$  is 0.2. The multiplicative margin and scale factor of CosFace loss [69] are set to 0.35 and 64, respectively. All images are aligned to  $112 \times 112$ , with five facial landmarks detected by MTCNN [70], and linearly normalized to  $[-1, 1]$ . For FT-Sel, SCAF is employed to synthesize the child faces, leading to a total of 60k high-quality synthesized faces. Afterward, the model is fine-tuned with a learning rate of 0.01, 10k iterations, and CosFace loss.

### 4.3 Evaluation on AIFR

Next, we evaluate MTLFace on several benchmark cross-age datasets, including CACD-VS [64], CALFW [62], AgeDB [63], FG-NET [65], and the proposed ECAF, to compare with the state-of-the-art methods. Note that MORPH is excluded since the version in [16], [22], [26] is prepared for commercial use only.

#### 4.3.1 Result on AgeDB Dataset

AgeDB [63] contains 16,488 face images of 568 distinct subjects with manually annotated age labels, which has four age-invariant face verification protocols under the different age gaps of face pairs: 5, 10, 20, and 30 years. Similar to the LFW [61], AgeDB is split into 10 folds for each protocol, where each fold consists of 300 intra-class and 300 inter-class pairs. We strictly follow the protocol of 30 years to perform 10-fold cross-validation since the protocol of 30 years is the most challenging one. We use models trained on SCAF to evaluate the performance on AgeDB for fair comparison. Table 2 shows the comparison results in terms of verification accuracy, demonstrating the superior performance of MTLFace over state-of-the-art methods.

#### 4.3.2 Result on CALFW Dataset

Cross-age LFW (CALFW) dataset [62] is designed for unconstrained face verification with large age gaps, which contains 12,176 face images of 4,025 individuals collected using the same identities in LFW. Similarly, we follow the same protocol as the LFW, where each fold consists of 600

TABLE 3  
Verification Rate (%) on CALFW Dataset for AIFR

Method	Accuracy (%)
HUMAN-Individual	82.32
HUMAN-Fusion	86.50
Center Loss [4]	85.48
SphereFace [73]	90.30
VGGFace2 [74]	90.57
ArcFace [51]	95.45
MTLFace (ours)	<b>95.98</b>

positive and negative pairs. We train the model on LCAF to evaluate our method on this dataset, and the results are shown in Table 3. Particularly, our method outperforms the recent state-of-the-art AIFR methods, establishing a new state-of-the-art on CALFW.

#### 4.3.3 Result on CACD-VS Dataset

The cross-age celebrity dataset (CACD) contains 163,446 face images of 2,000 celebrities in the wild, with significant variations in age, illumination, pose, and so on. Since collected by search engine, CACD is noisy with mislabeled and duplicate images. Therefore, a carefully annotated version, the CACD verification subset or CACD-VS [64], is constructed for fair comparison, which also follows the protocol of LFW. Table 4 presents the comparison of the proposed method with other state-of-the-art methods on CACD-VS. Our MTLFace surpasses the other state-of-the-art methods by a large margin, making an improvement of 0.18 against the second best.

#### 4.3.4 Result on FG-NET Dataset

FG-NET [65] is the most popular and challenging age dataset for AIFR, which consists of 1,002 face images from 82 subjects collected from the wild with huge age variations ranging from child to elder. We strictly follow the evaluation pipeline in [16], [22]. Specifically, the model is trained on SCAF and tested under the protocols of leave-one-out and MegaFace challenge 1 (MF1). In the leave-one-out protocol, faces are used to match the remaining faces, repeating 1,002 times. Table 5 reports the rank-1 recognition rate. Our method outperforms prior work

TABLE 4  
Verification Rate (%) on CACD-VS Dataset for AIFR

Method	Accuracy (%)
HFA [19]	84.40
CARC [64]	87.60
VGGFace [72]	96.00
Center Loss [4]	97.48
LF-CNN [20]	98.50
Marginal Loss [75]	98.95
OE-CNN [16]	99.20
AIM [26]	99.38
DAL [22]	99.40
MTLFace (ours)	<b>99.58</b>

TABLE 5  
Rank-1 Identification Rate (%) on FG-NET  
(Leave-one-out) Dataset for AIFR

Method	Rank-1 (%)
Park et al. [8]	37.40
Li et al. [76]	47.50
HFA [19]	69.00
MEFA [77]	76.20
CAN [78]	86.50
LF-CNN [20]	88.10
AIM [26]	93.20
DAL [22]	94.50
<b>MTLFace (ours)</b>	<b>95.00</b>

by a large margin. On the other hand, the MF1 contains an additional 1M images as the distractors in the gallery set from 690k different individuals, where models are evaluated under the large and small training set protocols. The small protocol requires the training set to be less than 0.5M images; this is strictly followed to evaluate our trained model on FG-NET, and the experimental results are reported in Table 6. Our method achieves competitive performance against other methods since the distractors in MF1 contain a large number of mislabeled probe and gallery face images.

#### 4.3.5 Result on ECAF Dataset

ECAF is the proposed cross-age benchmark for AIFR, which contains 5,265 faces including paired child and adult faces from 613 subjects. It follows the protocols of LFW and contains subset identities of LFW, which aims to evaluate the performance of face verification based on two tasks, i.e. <Adult, Child> and <Child, Child> with 6k and 4k pairs, respectively.

To establish baseline results on ECAF, we first evaluate humans on the task of distinguishing two faces of child and adult from the same person, following the procedures of LFW [61] and CACD-VS [64]. Specifically, we ask the users at Amazon Mechanical Turk (AMT) [81] to determine whether a given pair of faces belongs to the same person, and provide their confidence in the choices. To obtain more confident results, the users must have at least 95% approval rate at AMT, and each pair of faces has 10 results from different users. We report the human performance in Table 7. The voting performance significantly surpasses the average performance as the voting procedure ensembles the results from multiple users.

TABLE 6  
Rank-1 Identification Rate (%) on FG-NET  
(MF1) Dataset for AIFR

Method	Rank-1 (%)
FUDAN-CS_SDS [79]	25.56
SphereFace [73]	47.55
TNVP [32]	47.72
OE-CNN [16]	52.67
DAL [22]	57.92
<b>MTLFace (ours)</b>	<b>57.78</b>

TABLE 7  
Verification Rate (%) on ECAF Dataset for AIFR

Method	< Adult, Child >	< Child, Child >
Human, Average	73.34	68.62
Human, Voting	85.95	78.75
Softmax	85.03	88.25
CosFace [69]	85.72	90.75
ArcFace [51]	86.52	90.65
CurricularFace [80]	84.78	90.80
<b>MTLFace (ours)</b>	<b>87.55</b>	<b>91.20</b>

We then reproduced several baseline face recognition methods including Softmax, CosFace [69], ArcFace [51], and CurricularFace [80], and report the results of our MTLFace. Not surprisingly, MTLFace outperforms humans and other baseline face recognition methods. It is interesting to observe that the performance of humans on < Adult, Child > and < Child, Child > has the opposite results to those of CNN-based methods, i.e., the results of < Adult, Child > are better than the ones of < Child, Child >. A possible reason is that humans are not so confident in the choices of negative pairs in < Child, Child > compared to < Adult, Child >, as shown in Fig. 6. To obtain more insights into the AIFR task on ECAF, in supplemental Fig. 4, we showcase some example pairs with their cosine similarities on ECAF for two sub-tasks.

#### 4.3.6 Ablation Study

To investigate the efficacy of different modules in MTLFace, we perform ablation studies based on four benchmark datasets for AIFR by considering the following variants of our method. 1) Baseline: we remove all extra components other than the CosFace loss to train the face recognition model. 2) +Age: this variant is jointly trained under the supervision of both CosFace and age estimation loss, similar to [21], [22]. 3) +AFD (CA), +AFD (SA), +AFD (CBAM), +AFD: these four variants utilize the proposed attention-based feature decomposition to highlight the age-related information at different levels with different attention modules: CA [41], SA [42], CBAM [42], and the proposed one. 4) MTLFace w/ o FT: our proposed MTLFace is trained simultaneously by the AFD and cross-age domain adaptation loss without fine-tuning on the synthesized faces, which is the same as our preliminary version in [18]. 5) MTLFace w/ FT-All: directly using all the synthesized faces for fine-tuning. 6)




Fig. 6. Visualization of the confidence scores of humans regarding their choices, where a higher value indicates greater confidence.

TABLE 8  
Ablation Study on the Components of MTLFace for AIFR

Model	AgeDB-30	CALFW	CACD-VS	FG-NET
Baseline	95.52	94.27	99.12	93.64
+Age	95.32	94.35	99.15	93.88
+AFD (CA)	95.63	94.50	99.32	94.05
+AFD (SA)	95.85	94.43	99.25	94.38
+AFD (CBAM)	96.08	94.32	99.18	94.36
+AFD	95.90	94.48	99.30	94.58
MTLFace w/o FT	96.23	94.72	99.38	94.78
MTLFace w/ FT-All	95.88	94.45	99.22	93.98
MTLFace w/ FT-Sel	<b>96.45</b>	<b>94.97</b>	<b>99.45</b>	<b>95.00</b>

MTLFace w/ FT-Sel: using FT-Sel to select high-quality synthesized faces for fine-tuning.

Table 8 presents the experimental results. Note that the verification rate of the baseline model on AgeDB-30 is higher than those of ArcFace and DAAE since our training data are age-balanced, which is an important feature of our collected dataset. Even though the age estimation task is performed in the face recognition model, it cannot introduce any improvement in AIFR compared to the baseline model. On the other hand, AFD achieves a remarkable performance improvement on all cross-age datasets. Nevertheless, as AFD highlights the age-related information at both the channel and spatial levels in parallel, our method achieves consistent performance improvements, demonstrating its effectiveness compared to the single level (CA and SA) or sequential level (CBAM). Furthermore, the use of cross-age domain adversarial training leads to an additional performance improvement. At last, when all synthesized child faces are employed as the additional training data (i.e. MTLFace w/ FT-All), the performance drops as 1) the artifacts in low-quality synthesized faces would degrade the AIFR model, which has also been reported in [12]; and 2) using all synthesized children/ faces make the face recognition model focus more on children, where there are most adult faces in the testing sets. However, applying FT-Sel to AIFR (i.e. MTLFace w/ FT-Sel) has consistently improved the performance on several benchmark datasets. In supplemental Table 1, we show that FT-Sel has significantly improved the performance on children under 10 years old  $\langle$  Child, Child  $\rangle$ , and across ages  $\langle$  Adult, Child  $\rangle$  on the ECAF over our preliminary version [18] without synthesized faces.

#### 4.4 Evaluation on GFR

To validate the generalization ability of our MTLFace for GFR, we further conduct experiments on the LFW [61] and MegaFace Challenge 1 Facescrub (MF1-Facescrub) [82] datasets. LFW [61] is the most popular public benchmark dataset for GFR, which contains 13,233 face images from 5,749 subjects. MF1-Facescrub [82] uses the Facescrub dataset [83] of 106,863 face images from 530 celebrities as a probe set. The most challenging problem of MF1 is that it uses an additional 1M face images in the gallery set as distractors in face matching. That is, the results on MF1 are not as reliable as those on LFW due to the extremely noisy distractors in MF1. We strictly follow the same procedure as [16], [22]; i.e., the training dataset contains 0.5M images (SCAF).

TABLE 9  
General Face Recognition on LFW and MF1-Facescrub Dataset

Method	LFW	MF1-Facescrub
SphereFace [73]	99.42	72.73
CosFace [69]	99.33	77.11
OE-CNN [16]	99.35	N/A
DAL [22]	<u>99.47</u>	<b>77.58</b>
MTLFace (ours)	<b>99.55</b>	<u>77.33</u>

Table 9 reports the verification rate on LFW and the rank-1 identification rate on MF1-Facescrub against the state-of-the-art GFR methods. Our method achieves competitive performance on both datasets, demonstrating the strong generalization ability of our MTLFace. We highlight that our MTLFace can provide photorealistic synthesized faces to improve model interpretability, which is absent in other methods [16], [22].

#### 4.5 Evaluation on FAS

We further evaluate the model trained on SCAF for FAS.

##### 4.5.1 Qualitative Results

Fig. 7 presents some sample results on the external datasets including MORPH and FG-NET. Our method is able to simulate the face age synthesis process between age groups with high visual fidelity. Although variations exist in terms of race, gender, expression, and occlusion, the synthesized faces are still photorealistic, with natural details in the skin, muscles, and wrinkles while consistently preserving identities, thus confirming the generalization ability of the proposed method.

##### 4.5.2 Comparisons With Prior Work

We also conduct qualitative comparisons with prior work including CAEE [14], AIM [26], PAG-GAN [13], IPCGAN [12], DAAE [37], and PFA-GAN [39] on MORPH, FG-NET, and CACD datasets. Fig. 8 shows that both CAEE and AIM produce oversmoothed faces due to their image reconstruction and other superior methods age faces with more photorealistic aging effects but subtle changes. However, our MTLFace can synthesize faces with facial shape changes and stronger aging effects even though the age gaps are large (see the comparisons with DAAE). In addition to the image quality, MTLFace has not been re-trained on these three benchmark datasets and owns finer age split (7 versus 4 age groups in most previous literature), which demonstrates the strong generalization ability of MTLFace against other competitors. Note that the results of competitors are directly referred from their own papers for a fair comparison, which is widely adopted in the FAS literature such as [10], [11], [13], [37], [59] to avoid any bias or error caused by self-implementation.

##### 4.5.3 Quantitative Comparisons

We quantitatively evaluate MTLFace over previous competitive methods including CAEE [14], IPCGAN [12], and S2GAN [59], by the following three metrics:

- 1) age accuracy: we trained a ResNet-100 model on 80% faces of LCAF using  $\ell_{AE}$  as the loss function to predict the ages of all synthesized faces, where the proportion

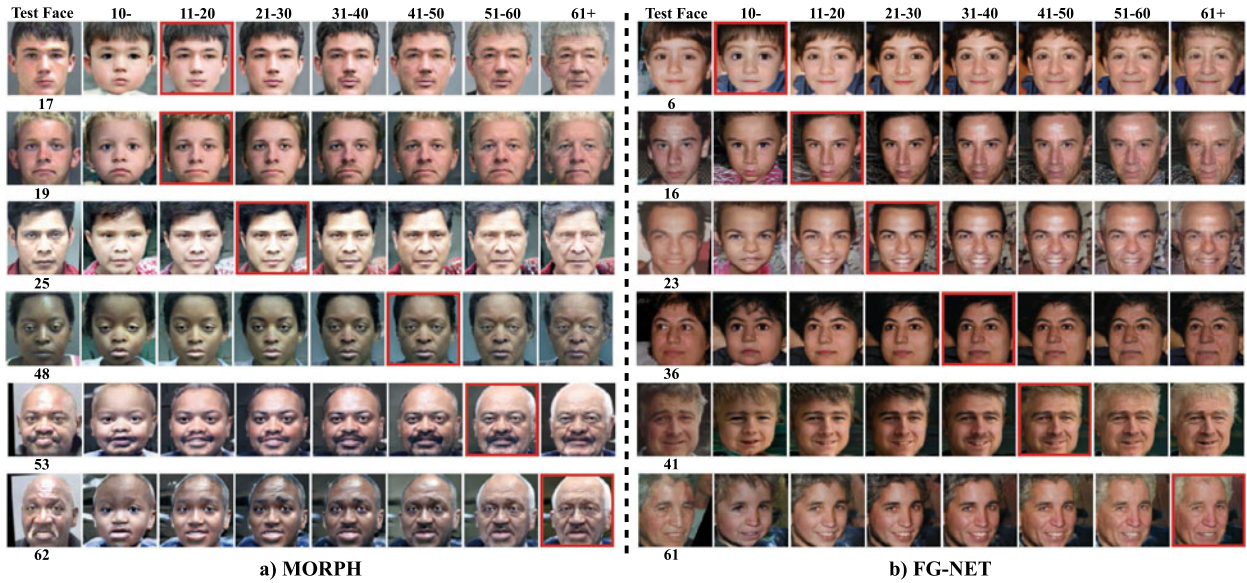


Fig. 7. Qualitative results by applying our MTLFace trained on the SCAF dataset to two external datasets : a) MORPH; and b) FG-NET. Red boxes indicate input faces.

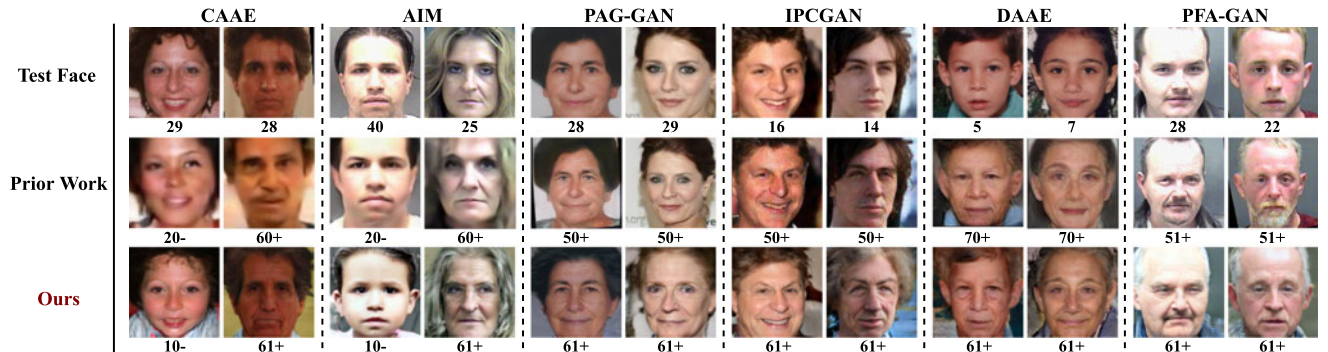


Fig. 8. Comparisons with prior work on FG-NET, MORPH, and CACD datasets. We show the test faces in the first row with the real age labels below each image. The second row presents two sample results of prior work, with the target age below each image. The third row shows our results of the same input faces for the same face aging and rejuvenation as in prior work. Zoom in for a better view of image details.

TABLE 10

Quantitative Comparisons Between our MTLFace and the State-of-the-art Face Aging/Rejuvenation Methods in the Form of  $a/b/c$ , Where  $a$ ,  $b$ , and  $c$  Represent the Mean Values of age Accuracy (%), Mean Absolute Error, and Identity Preservation (Cosine Similarity) Computed Over all age Mappings, Respectively

Method	MORPH	FG-NET	CACD	ECAF
CAAE [14]	39.77/8.83/0.131±0.089	40.72/9.32/0.146±0.097	39.50/8.74/0.118±0.092	41.36/8.75/0.139±0.100
IPCGAN [12]	58.44/4.95/0.608±0.109	61.22/4.32/0.453±0.122	62.04/3.99/0.452±0.129	61.92/4.16/0.511±0.133
S2GAN [59]	53.85/5.47/0.369±0.105	60.26/4.37/0.293±0.111	60.19/4.31/0.270±0.110	62.35/3.98/0.297±0.117
MTLFace (ours)	<b>67.07/2.98/0.652±0.091</b>	<b>71.33/3.10/0.648±0.094</b>	<b>71.94/2.55/0.630±0.096</b>	<b>72.14/2.28/0.620±0.096</b>
w/o AIFR	51.79/5.49/0.219±0.095	48.60/7.68/0.235±0.093	50.69/6.50/0.197±0.095	52.31/6.14/0.217±0.103
w/o Multi-level ICMs	64.80/3.35/0.606±0.090	69.35/3.27/0.628±0.094	67.79/3.02/0.613±0.098	70.77/2.65/0.601±0.099
w/o ICM	62.60/3.86/0.605±0.072	65.46/3.65/0.612±0.080	62.55/3.87/0.591±0.099	67.34/3.16/0.584±0.088

of the predicted ages falling into the target age groups is the age accuracy;

- 2) mean absolute error (MAE): MAE between predicted and ground-truth ages used in [49], [84], [85] is adopted as one of evaluation metric. Here, the mean ages of target age groups is the ground-truth age label for synthesized faces; e.g., 5 years old for 10-age group, similar to [86]; and
- 3) identity preservation: an external well-trained face recognition model, the ResNet-100 network pre-

trained on the MS-Celeb-1M dataset from [51], is used for fair comparisons to compute the cosine similarity between the input and synthesized faces.

We reproduce CAAE, IPCGAN, and S2GAN on the SCAF dataset for fair comparisons, and then directly apply them to three external cross-age datasets: MORPH [87], FG-NET [65], CACD [64], and the proposed ECAF. Table 10 presents the quantitative results of different face aging/rejuvenation methods, including the competitive methods, our proposed MTLFace and its two variants (w/o multi-

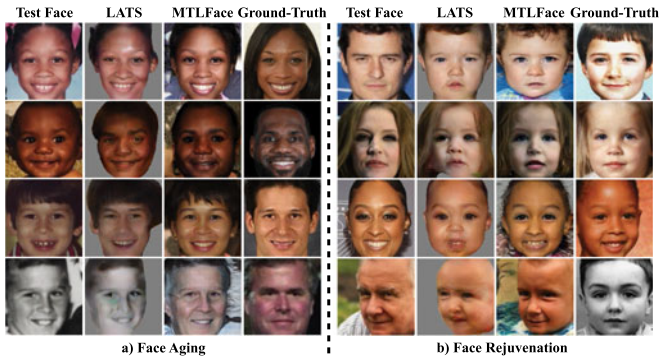


Fig. 9. Comparison with ground-truth images on ECAF. We synthesize the faces by a) face aging and b) face rejuvenation to the same age group.

level ICMs and w/o ICM), in terms of three evaluation metrics. MTLFace outperforms other competitors by a clear margin; this is a direct result of AIFR and ICM with multi-level architecture. Without ICM, MTLFace reduces to a common cGANs-based method that uses one-hot encoding to control face aging/rejuvenation at the group level. Remarkably, the MTLFace without ICM still outperforms these two baseline methods, implying that our multi-learning framework with attention-based feature decomposition is effective in improving the quantitative results. We also ablate AIFR from MTLFace for FAS task; i.e., the encoder was trained with FAS without the training signals of AIFR. Without AIFR, FAS cannot produce faces with the same identity at the target age group as expected, which shows the essential importance of the AIFR task for FAS; see supplemental Fig. 3 for qualitative comparisons between MTLFace with and without AIFR.

#### 4.5.4 Comparison With Ground-Truth Faces

ECAF provides pairs of child and adult faces from the same person, which enables a comparison between the synthesized faces and ground-truth faces. To this end, we also conduct comparisons with the trained CAEE, IPCGAN, S2GAN, and one of the current state-of-the-art methods LATS [38]. Note that we use the public pre-trained model of LATS to produce the synthesized faces, which are then aligned with MTCNN [70] to avoid potential self-implementation bias.

Fig. 9 visualizes the aging/rejuvenation process for MTLFace and LATS with the reference of ground-truth faces. Even though there are large age gaps from child to adult, MTLFace can synthesize realistic faces with similar aging/rejuvenation effects compared to ground-truth faces in terms of identities. On the contrary, LATS produces faces with severe artifacts for these low-quality input faces—low resolution and extreme poses—which indicates that MTLFace has a stronger generalization ability.

To further demonstrate the effects of face age synthesis, we quantitatively evaluate the synthesized faces against corresponding ground-truth adult/child faces in terms of cosine similarity. First, the positive pairs from ECAF are aged/rejuvenated to the same age groups as the ground-truth faces. Then, the trained face recognition model of Cos-Face is employed for fair comparisons to compute the cosine similarity between the synthesized and ground-truth faces.

TABLE 11  
Cosine Similarity on Ground-Truth Face Pairs of ECAF for FAS

Method	Aging	Rejuvenation
Ground-Truth	$0.182 \pm 0.111$	$0.182 \pm 0.111$
CAAE [14]	$0.030 \pm 0.077$	$0.078 \pm 0.088$
IPCGAN [12]	$0.137 \pm 0.099$	$0.193 \pm 0.098$
S2GAN [59]	$0.065 \pm 0.090$	$0.135 \pm 0.095$
LATS [38]	$0.103 \pm 0.090$	$0.172 \pm 0.106$
MTLFace (ours)	<b><math>0.164 \pm 0.106</math></b>	<b><math>0.210 \pm 0.112</math></b>

The higher values, the better identity preservation in synthesized faces.

The results are reported in Table 11 in terms of face aging and rejuvenation. Compared to the results of the ground-truth, the face rejuvenation of MTLFace significantly improves the discrimination of the face recognition model (see increased cosine similarity). However, it drops slightly for face aging. Similarly, this phenomenon turns out consistently on other methods for face aging. We think the reason is that face aging is much more difficult than face rejuvenation, so the unexpected ghost artifacts harm the recognition performance. Although face aging is largely influenced by individuals, health conditions, personal habits, and so on, MTLFace still performs better in preserving personal identities than other methods in terms of ground-truth, demonstrating the effectiveness of AIFR task.

#### 4.5.5 Continuous Face Age Synthesis

Thanks to the StyleGAN-based architecture used in the face age synthesis task, MTLFace can interpolate the missing classes in the latent space between two adjacent age groups to achieve continuous age progression and regression. Fig. 10 shows examples of continuous face age synthesis for MTLFace. Compared to the synthesis only on discrete age groups in previous works [10], [11], [13], [37], MTLFace can directly interpolate in the latent space to synthesize realistic faces with smooth aging sequences. Note that the continuous synthesis is *jointly* achieved by the StyleGAN architecture and the proposed ICM, whereas only StyleGAN cannot perform well on FAS. In supplementary Section 2, we clarify the differences between StyleGAN and our MTLFace for continuous synthesis, summarize the advantages of MTLFace, and conduct extensive quantitative and qualitative comparisons.

#### 4.5.6 Training Visualization

This section examines the training stability of MTLFace. That is, when AIFR is not well trained at the early stage of training, FAS may not preserve the identities of input faces, which consequently harms the quality of synthesized faces. However, we argue that the unstable training would not arise in MTLFace due to the following reasons:

- 1) In addition to the identity preservation loss, another learned perceptual image patch similarity (LPIPS) loss [58] is employed to further encourage the perceptual consistency between input and synthesized faces. Although AIFR cannot provide FAS with a good training signal at the early stage of training, the LPIPS loss can still encourage FAS to reconstruct the input faces.

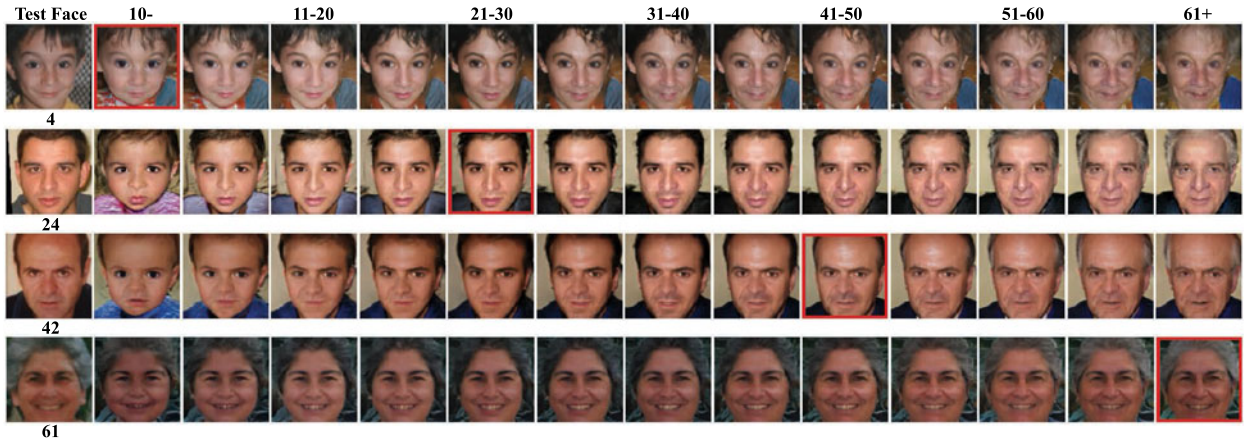


Fig. 10. Sample results of continuous face age progression and regression for MTLFace on the FG-NET dataset. We show the test faces in the first column, with the age below the image. The red boxes indicate the reconstructed faces in the same age groups as the input faces. We synthesize the continuous aging/rejuvenation faces by interpolating between the two latent variables of adjacent age groups. Zoom in for better view.

- 2) AIFR can extract discriminative representations that are good enough for training FAS at the beginning of training.

To validate these two points, we visualize the training process of both AIFR and FAS in Fig. 11. At the first 1k iterations, AIFR can achieve a good performance on LFW verification while FAS cannot reconstruct the same identities. The results indicate that although AIFR has not reached its best performance, it can still provide good enough training signals for FAS. In terms of synthesized faces, FAS can still reconstruct pleasing faces benefitting from the LPIPS loss, and become better at preserving the identities with the training of AIFR.

## 5 DISCUSSION

Although the proposed MTLFace has achieved state-of-the-art performance for the AIFR task and can produce pleasing synthesized faces for the FAS task, we acknowledge some limitations in this work. Specifically, although StyleGAN-based decoder has significantly improved the image quality

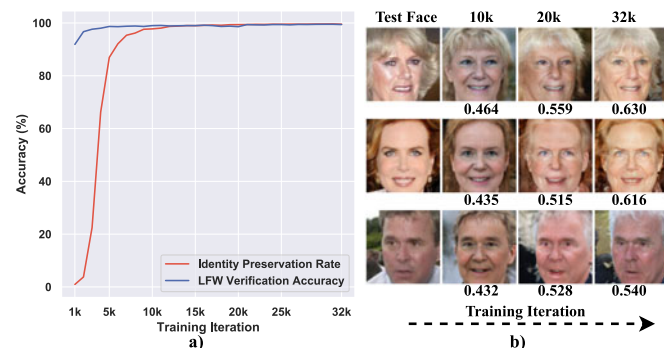


Fig. 11. Visualization of the training process in terms of qualitative and quantitative results. a) Identity preservation rate of face age synthesis and LFW verification accuracy of face recognition across training. The identity preservation rate denotes the ratio of the pairs whose cosine similarity is larger than the threshold computed on LFW. Even at the early stage of the training process, AIFR performs well so that it can provide good enough training signals to encourage FAS for photo-realistic faces. b) Synthesized faces of the 61+ age group at different training iterations with the cosine similarities between input and synthesized faces shown below. FAS can reconstruct pleasing faces at the early stage of training, and become better at preserving identities with the training of AIFR.

of synthesized faces and achieved continuous face age synthesis, we found that it cannot well preserve the background of input faces. Since face age synthesis mainly focuses on facial changes, the background may be relatively less important. This may be addressed by the following two solutions. First solution is to use another pixel-wise loss such as mean absolute loss, which may inevitably reduce the aging/rejuvenation effects. Second solution is to extract the face area from the input images prior to being fed into the synthesis network to reduce the influence of background like LATS [38], which may, however, introduce extra computational cost for face segmentation.

## 6 CONCLUSION

In this paper, we proposed a multi-task learning framework, termed MTLFace, to achieve AIFR and FAS simultaneously. We proposed two novel modules: AFD to decompose the features into age- and identity-related features, and ICM to achieve identity-level FAS. We also proposed a novel selective fine-tune strategy to boost AIFR that selects and leverages the high-quality synthesized faces. Extensive experiments on both cross-age and general benchmark datasets for face recognition demonstrate the superiority of our MTLFace. With multi-level ICMs, MTLFace can be significantly better at preserving the identities of input faces by multi-level skip connections, and improve the age accuracy due to the multi-level age conditions. The newly collected large-scale cross-age training dataset and benchmark could further advance the development of AIFR and FAS.

## ACKNOWLEDGMENTS

The authors would like to thank the associate editor and two anonymous reviewers for their valuable comments, which greatly improved the quality of this article.

## REFERENCES

- [1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 770–778.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 84–90, 2017.

- [3] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Representation.*, 2015, pp. 1–14.
- [4] Y. Wen, K. Zhang, Z. Li, and Y. Qiao, "A discriminative feature learning approach for deep face recognition," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 499–515.
- [5] Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "MagFace: A universal representation for face recognition and quality assessment," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 14 225–14 234.
- [6] X. Geng, Z.-H. Zhou, and K. Smith-Miles, "Automatic age estimation based on facial aging patterns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 12, pp. 2234–2240, Dec. 2007.
- [7] A. Lanitis, C. J. Taylor, and T. F. Cootes, "Toward automatic simulation of aging effects on face images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 4, pp. 442–455, Apr. 2002.
- [8] U. Park, Y. Tong, and A. K. Jain, "Age-invariant face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 5, pp. 947–954, May 2010.
- [9] I. Goodfellow et al., "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [10] Q. Li, Y. Liu, and Z. Sun, "Age progression and regression with spatial attention modules," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 11378–11385.
- [11] Y. Liu, Q. Li, and Z. Sun, "Attribute-aware face aging with wavelet-based generative adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 11 877–11 886.
- [12] Z. Wang, X. Tang, W. Luo, and S. Gao, "Face aging with identity-preserved conditional generative adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7939–7947.
- [13] H. Yang, D. Huang, Y. Wang, and A. K. Jain, "Learning face age progression: A pyramid architecture of GANs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 31–39.
- [14] Z. Zhang, Y. Song, and H. Qi, "Age progression/regression by conditional adversarial autoencoder," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5810–5818.
- [15] D. Deb, D. Aggarwal, and A. K. Jain, "Finding missing children: Aging deep face features," 2019, *arXiv: 1911.07538*.
- [16] Y. Wang et al., "Orthogonal deep features decomposition for age-invariant face recognition," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 738–753.
- [17] Y. Ganin et al., "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [18] Z. Huang, J. Zhang, and H. Shan, "When age-invariant face recognition meets face age synthesis: A multi-task learning framework," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 7282–7291.
- [19] D. Gong, Z. Li, D. Lin, J. Liu, and X. Tang, "Hidden factor analysis for age invariant face recognition," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013, pp. 2872–2879.
- [20] Y. Wen, Z. Li, and Y. Qiao, "Latent factor guided convolutional neural networks for age-invariant face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4893–4901.
- [21] T. Zheng, W. Deng, and J. Hu, "Age estimation guided convolutional neural network for age-invariant face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2017, pp. 1–9.
- [22] H. Wang, D. Gong, Z. Li, and W. Liu, "Decorrelated adversarial learning for age-invariant face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3527–3536.
- [23] J. Lee, J. Yun, S. Park, Y. Kim, and J. Choo, "Improving face recognition with large age gaps by learning to distinguish children," in *Proc. Brit. Mach. Vis. Conf.*, 2021.
- [24] X. Hou, Y. Li, and S. Wang, "Disentangled representation for age-invariant face recognition: A mutual information minimization perspective," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 3692–3701.
- [25] J.-C. Xie, C.-M. Pun, and K.-M. Lam, "Implicit and explicit feature purification for age-invariant facial representation learning," *IEEE Trans. Inf. Forensics Secur.*, vol. 17, pp. 399–412, 2022.
- [26] J. Zhao et al., "Look across elapse: Disentangled representation learning and photorealistic cross-age face synthesis for age-invariant face recognition," in *Proc. 33rd AAAI Conf. Artif. Intell. 31st Innov. Appl. Artif. Intell. Conf. 9th AAAI Symp. Educ. Adv. Artif. Intell.*, 2019, pp. 9251–9258.
- [27] N. Ramanathan and R. Chellappa, "Modeling age progression in young faces," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2006, pp. 387–394.
- [28] N. Ramanathan and R. Chellappa, "Modeling shape and textural variations in aging faces," in *Proc. IEEE Conf. Autom. Face Gesture Recognit.*, 2008, pp. 1–8.
- [29] J. Suo, X. Chen, S. Shan, W. Gao, and Q. Dai, "A concatenational graph evolution aging model," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2083–2096, Nov. 2012.
- [30] I. Kemelmacher-Shlizerman, S. Suwajanakorn, and S. M. Seitz, "Illumination-aware age progression," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 3334–3341.
- [31] D. A. Rowland and D. I. Perrett, "Manipulating facial appearance through shape and color," *IEEE Comput. Graph. Appl.*, vol. 15, no. 5, pp. 70–76, Sep. 1995.
- [32] C. N. Duong, K. Gia Quach, K. Luu, N. Le, and M. Savvides, "Temporal non-volume preserving approach to facial age-progression and age-invariant face recognition," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 3735–3743.
- [33] W. Wang et al., "Recurrent face aging," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2378–2386.
- [34] M. Mirza and S. Osindero, "Conditional generative adversarial nets," 2014, *arXiv:1411.1784*.
- [35] J. Despois, F. Flament, and M. Perrot, "Agingmapgan (AMGAN): High-resolution controllable face aging with spatially-aware conditional GANs," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 613–628.
- [36] Z. Huang, S. Chen, J. Zhang, and H. Shan, "AgeFlow: Conditional age progression and regression with normalizing flows," in *Proc. Int. Joint Conf. Artif. Intell.*, 2021.
- [37] P. Li, H. Huang, Y. Hu, X. Wu, R. He, and Z. Sun, "Hierarchical face aging through disentangled latent characteristics," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 86–101.
- [38] R. Or-El, S. Sengupta, O. Fried, E. Shechtman, and I. Kemelmacher-Shlizerman, "Lifespan age transformation synthesis," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 739–755.
- [39] H. Huang, S. Chen, J. Zhang, and H. Shan, "PFA-GAN: Progressive face aging with generative adversarial network," *IEEE Trans. Inf. Forensics Secur.*, vol. 16, pp. 2031–2045, 2021.
- [40] H. Yang, D. Huang, Y. Wang, and A. K. Jain, "Learning continuous face age progression: A pyramid of GANs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 2, pp. 499–515, Feb. 2021.
- [41] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7132–7141.
- [42] S. Woo, J. Park, J.-Y. Lee, and I. So Kweon, "CBAM: Convolutional block attention module," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 3–19.
- [43] Y. Zhang, B. Kang, B. Hooi, S. Yan, and J. Feng, "Deep long-tailed learning: A survey," 2021, *arXiv:2110.04596*.
- [44] M. Hong, J. Choi, and G. Kim, "StyleMix: Separating content and style for enhanced data augmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 14857–14865.
- [45] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2414–2423.
- [46] H. Qiu, B. Yu, D. Gong, Z. Li, W. Liu, and D. Tao, "SynFace: Face recognition with synthetic data," in *Proc. Int. Conf. Comput. Vis.*, 2021, pp. 10880–10890.
- [47] S. Zhang, Z. Li, S. Yan, X. He, and J. Sun, "Distribution alignment: A unified framework for long-tail visual recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 2361–2370.
- [48] R. Rothe, R. Timofte, and L. Van Gool, "DEX: Deep expectation of apparent age from a single image," in *Proc. IEEE Int. Conf. Comput. Vis. Workshop*, 2015, pp. 10–15.
- [49] R. Rothe, R. Timofte, and L. V. Gool, "Deep expectation of real and apparent age from a single image without facial landmarks," *Int. J. Comput. Vis.*, vol. 126, no. 2, pp. 144–157, 2018.
- [50] H. Wang, H. He, and D. Katabi, "Continuously indexed domain adaptation," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 9898–9907.
- [51] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4690–4699.
- [52] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4401–4410.
- [53] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8110–8119.
- [54] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2117–2125.


- [55] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 1501–1510.
- [56] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1125–1134.
- [57] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2794–2802.
- [58] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 586–595.
- [59] Z. He, M. Kan, S. Shan, and X. Chen, "S2GAN: Share aging factors across ages and share aging trends among individuals," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 9440–9449.
- [60] M. Azure, "Microsoft azure cognitive services facial recognition," 2022. [Online]. Available: <https://azure.microsoft.com/services/cognitive-services/face/>
- [61] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," University of Massachusetts, Amherst, Tech. Rep. 07–49, Oct. 2007.
- [62] T. Zheng, W. Deng, and J. Hu, "Cross-age LFW: A database for studying cross-age face recognition in unconstrained environments," 2017, *arXiv:1708.08197*.
- [63] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotsia, and S. Zafeiriou, "AgeDB: The first manually collected, in-the-wild age database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Worksh.*, 2017, pp. 51–59.
- [64] B.-C. Chen, C.-S. Chen, and W. H. Hsu, "Face recognition and retrieval using cross-age reference coding with cross-age celebrity dataset," *IEEE Trans. Multimedia*, vol. 17, no. 6, pp. 804–815, Jun. 2015.
- [65] FG-NET aging database, 2014. [Online]. Available: [https://yanweifu.github.io/FG\\_NET\\_data](https://yanweifu.github.io/FG_NET_data)
- [66] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," 2016, *arXiv:1607.08022*.
- [67] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral normalization for generative adversarial networks," in *Proc. Int. Conf. Learn Represent.*, 2018, pp. 1–26.
- [68] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [69] H. Wang et al., "CosFace: Large margin cosine loss for deep face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 5265–5274.
- [70] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Process. Lett.*, vol. 23, no. 10, pp. 1499–1503, Oct. 2016.
- [71] C. Sagonas, E. Ververas, Y. Panagakis, and S. Zafeiriou, "Recovering joint and individual components in facial data," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 11, pp. 2668–2681, Nov. 2017.
- [72] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in *Proc. Brit. Mach. Vis. Conf.*, 2015, pp. 471–478.
- [73] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, "SphereFace: Deep hypersphere embedding for face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 212–220.
- [74] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A dataset for recognising faces across pose and age," in *Proc. IEEE Conf. Autom. Face Gesture Recognit.*, 2018, pp. 67–74.
- [75] J. Deng, Y. Zhou, and S. Zafeiriou, "Marginal loss for deep face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Worksh.*, 2017, pp. 60–68.
- [76] Z. Li, U. Park, and A. K. Jain, "A discriminative model for age invariant face recognition," *IEEE Trans. Inf. Forensics Secur.*, vol. 6, no. 3, pp. 1028–1037, Sep. 2011.
- [77] D. Gong, Z. Li, D. Tao, J. Liu, and X. Li, "A maximum entropy feature descriptor for age invariant face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 5289–5297.
- [78] C. Xu, Q. Liu, and M. Ye, "Age invariant face recognition and retrieval by coupled auto-encoder networks," *Neurocomputing*, vol. 222, pp. 62–71, 2017.
- [79] Z. Wang, K. He, Y. Fu, R. Feng, Y.-G. Jiang, and X. Xue, "Multi-task deep neural network for joint face recognition and facial attribute prediction," in *Proc. ACM Int. Conf. Multimedia Retrieval*, 2017, pp. 365–374.
- [80] Y. Huang et al., "CurricularFace: Adaptive curriculum learning loss for deep face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 5900–5909.
- [81] Amazon mechanical turk, 2021. [Online]. Available: <https://www.mturk.com/>
- [82] I. Kemelmacher-Shlizerman, S. M. Seitz, D. Miller, and E. Bros-sard, "The MegaFace benchmark: 1 million faces for recognition at scale," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4873–4882.
- [83] H.-W. Ng and S. Winkler, "A data-driven approach to cleaning large face datasets," in *Proc. IEEE Int. Conf. Image Process.*, 2014, pp. 343–347.
- [84] Z. Tan, J. Wan, Z. Lei, R. Zhi, G. Guo, and S. Z. Li, "Efficient group-n encoding and decoding for facial age estimation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 11, pp. 2610–2623, Nov. 2018.
- [85] C. Zhu, X. Li, J. Li, and S. Dai, "Improving robustness of facial landmark detection by defending against adversarial attacks," in *Proc. Int. Conf. Comput. Vis.*, 2021, pp. 11751–11760.
- [86] H. Zhu, Z. Huang, H. Shan, and J. Zhang, "Look globally, age locally: Face aging with an attention mechanism," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 1963–1967.
- [87] K. Ricanek and T. Tesafaye, "MORPH: A longitudinal image database of normal adult age-progression," in *Proc. IEEE 7th Int. Conf. Autom. Face Gesture Recognit.*, 2006, pp. 341–345.



**Zhizhong Huang** (Graduate Student Member, IEEE) received the BS degree from the Sichuan University, Chengdu, China, in 2019. He is currently working toward the Ph.D. degree with the School of Computer Science, Fudan University, Shanghai, China. His research interests include machine/deep learning, computer vision, face aging, and image translation.



**Junping Zhang** (Senior Member, IEEE) received the BS degree in automation from Xiangtan University, China, in 1992, the MS degree in control theory and control engineering from Hunan University, Changsha, China, in 2000, and the PhD degree in intelligent systems and pattern recognition from the Institution of Automation, Chinese Academy of Sciences, in 2003. He has been a professor with the School of Computer Science, Fudan University, since 2006. His research interests include machine learning, image processing, biometric authentication, and intelligent transportation systems. He has been an associate editor of *IEEE Intelligent System* since 2009 and was an associate editor of *IEEE Transactions on Intelligent Transportation Systems* (2010–2018).



**Hongming Shan** (Senior Member, IEEE) received the PhD degree in machine learning from Fudan University, in 2017. From 2017 to 2020, he was a postdoctoral research associate and Research Scientist with Rensselaer Polytechnic Institute, USA. He is currently an associate professor with the Institute of Science and Technology for Brain-inspired Intelligence, Fudan University, and also a "Qiusuo" Research Leader with the Shanghai Center for Brain Science and Brain-inspired Technology. His research focuses on developing machine learning algorithms for biomedical imaging. He was recognized with Youth Outstanding Paper Award with World Artificial Intelligence Conference, 2021.

► For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/csdl](http://www.computer.org/csdl).