

AgeFlow: Conditional Age Progression and Regression with Normalizing Flows

Zhizhong Huang¹, Shouzhen Chen¹, Junping Zhang¹ and Hongming Shan^{2,3*}

¹Shanghai Key Lab of Intelligent Information Processing, School of Computer Science
Fudan University, Shanghai 200433, China

²Institute of Science and Technology for Brain-inspired Intelligence and MOE Frontiers Center
for Brain Science, Fudan University, Shanghai 200433, China

³Shanghai Center for Brain Science and Brain-inspired Technology, Shanghai 200031, China
{zzhuang19, chensz19, jpzhang, hmshan}@fudan.edu.cn

Abstract

Age progression and regression aim to synthesize photorealistic appearance of a given face image with aging and rejuvenation effects, respectively. Existing generative adversarial networks (GANs) based methods suffer from the following three major issues: 1) unstable training introducing strong ghost artifacts in the generated faces, 2) unpaired training leading to unexpected changes in facial attributes such as genders and races, and 3) non-bijective age mappings increasing the uncertainty in the face transformation. To overcome these issues, this paper proposes a novel framework, termed AgeFlow, to integrate the advantages of both flow-based models and GANs. The proposed AgeFlow contains three parts: an encoder that maps a given face to a latent space through an invertible neural network, a novel invertible conditional translation module (ICTM) that translates the source latent vector to target one, and a decoder that reconstructs the generated face from the target latent vector using the same encoder network; all parts are invertible achieving bijective age mappings. The novelties of ICTM are two-fold. First, we propose an attribute-aware knowledge distillation to learn the manipulation direction of age progression while keeping other unrelated attributes unchanged, alleviating unexpected changes in facial attributes. Second, we propose to use GANs in the latent space to ensure the learned latent vector indistinguishable from the real ones, which is much easier than traditional use of GANs in the image domain. Experimental results demonstrate superior performance over existing GANs-based methods on two benchmarked datasets. The source code is available at <https://github.com/Hzzone/AgeFlow>.

1 Introduction

Age progression and regression, also known as face aging and rejuvenation, are to synthesize the appearance of a given face image at different ages with natural aging/rejuvenation

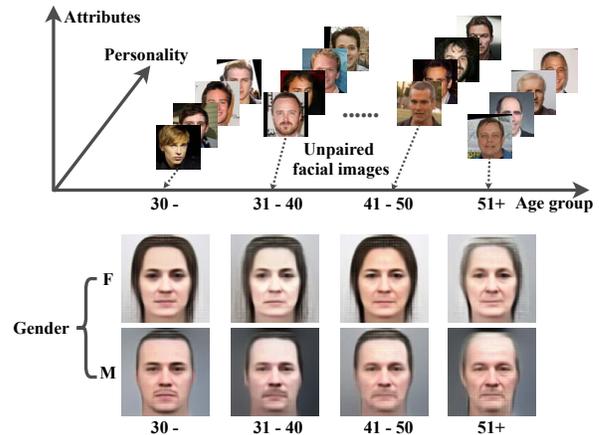


Figure 1: Top: Illustration of a face latent space learned by a flow-based generative model, in which faces align with different facial conditions. Bottom: The faces reconstructed from the *average* latent variables for each attribute pair, e.g. genders and age groups.

effects. Despite their broad applications ranging from entertainment to social security [Park *et al.*, 2010; Huang *et al.*, 2021c], the intrinsic complexities such as complicated translation patterns, various expressions, and diversified races pose challenges to existing methods. Recent studies [Li *et al.*, 2020a; Or-EI *et al.*, 2020; Huang *et al.*, 2021b] resort to generative adversarial networks (GANs) [Goodfellow *et al.*, 2014] and its variant conditional GANs (cGANs) [Mirza and Osindero, 2014] to take advantage of massive unpaired age data, achieving more impressive results than conventional methods [Wang *et al.*, 2016].

In spite of significant progress in meeting three important requirements of face aging/rejuvenation—image quality, age accuracy, and identity preservation, these GANs-based methods suffer the following shortcomings. **First**, the unstable training of GANs introduces strong ghost artifacts in the generated faces, which severely compromises the image quality of age progression/regression. **Second**, the unpaired training of GANs in the image domain may lead to unexpected changes in facial attributes due to the neglect of some important conditional information of the input (e.g., gender and race), thus challenging the identity preservation of input faces. **Third**, these GANs-based methods han-

*Corresponding author

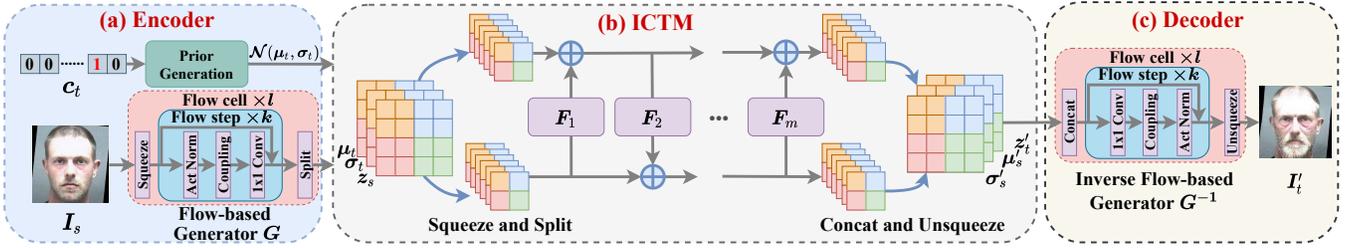


Figure 2: The overall framework of the proposed AgeFlow. (a) The encoder maps an input image I_s into z_s in the latent space using a flow-based generator G , and a prior generation module converts one-hot encoding of target age c_t to a Gaussian distribution, $\mathcal{N}(\mu_t, \sigma_t)$. (b) The invertible conditional translation module (ICTM) takes z_s , μ_t and σ_t as input and achieves conditional age progression/regression with normalizing flows in the latent space, outputting the progressed/regressed latent variables z'_t and source age distribution from I_s to keep the cycle consistency. ICTM contains m flows, each of which contains two convolutional layers, one channel attention module, and one convolutional layer initialized by zero. (c) The inverse function of G reconstructs the generated face I'_t from z'_t .

handle all age mappings together with one single network under the control of an age condition, increasing the uncertainty in the face transformation with unsatisfied age accuracy. Several methods have been proposed to address these problems. Initial attempts [Yang *et al.*, 2018; Liu *et al.*, 2019; Huang *et al.*, 2021a] mainly focus on training several independent models for each age mapping. Considering the different patterns of age progression and regression, [Song *et al.*, 2018] used two separate generators to model both age changing processes separately, and [Li *et al.*, 2020b] further extended it with a spatial attention mechanism. However, they usually suffer from considerably computational costs and are limited with the age label as prior.

Inspired by the great success of flow-based generative models [Dinh *et al.*, 2015; Dinh *et al.*, 2017; Kingma and Dhariwal, 2018] in exact latent-variable inference, this paper exploits the advantages of flow-based generative models and GANs for both age progression/regression. Most importantly, a more stable training process can be achieved by the flow-based generative models as they are directly optimized by the log-likelihood objective function instead of an adversarial competition. Figure 1 (top) illustrates that the faces in the latent space align with the facial attributes, personality, and ages. Therefore, the average latent variables for each attribute can be computed as the direction for manipulating the attribute-aware latent space. In this paper, we mainly focus on manipulating the ages of faces while keeping other attributes unchanged for the present age progression and regression tasks. Therefore, an average latent vector for each attribute pair can keep the attribute consistent during face aging/rejuvenation. Taking the genders as an example, Figure 1 (bottom) shows the reconstructed prototypes, *i.e.*, average latent vectors of different age groups and genders. Remarkably, the wrinkle appears, hair becomes white, and beard grows *only for male* on the prototype faces as expected when faces age from left to right. However, the original flow-based models lack the capability of conditionally facial manipulation. To achieve conditional age progression/regression similar to the cGANs-based method, the learned manipulation direction can be used to guide the generative models in a knowledge distilling manner with the advantage of GANs to ensure the learned latent vector indistinguishable from the real one.

Contribution. The contributions are summarized as follows. 1) We propose a novel framework, termed AgeFlow, to integrate the advantages of both flow-based models and GANs, which contains three parts—encoder, invertible conditional translation module (ICTM), and decoder; all are invertible reducing the uncertainty in face transformation. 2) We introduce an attribute-aware knowledge distillation to learn the manipulation direction of face progression while keeping other unrelated attributes unchanged, which can reduce unexpected changes in facial attributes. 3) We apply GANs in the latent space to ensure the learned latent vector indistinguishable from the real ones, which is much easier than the traditional use of GANs in the image domain.

2 Methodology

2.1 Problem Formulation

Given a face image I_s at source age c_s , age progression/regression aims to learn a neural network to naturally translate I_s into I_t conditioned on a given target age c_t . Formally, the age progression/regression process of the cGANs-based methods can be formulated as:

$$I'_t = G(I_s, c_t), \quad (1)$$

where I'_t represents the generated face from I_s . However, the translation $G : I_s \rightarrow I_t$ is highly under-constrained due to the unpaired training of GANs, misleading the model to learn the patterns other than face aging/rejuvenation. One possible solution is to regularize the mappings with cycle-consistency [Zhu *et al.*, 2017], which learns backwards translation $I_t \rightarrow I_s$ using a separate model [Song *et al.*, 2018; Li *et al.*, 2020b] to approximately reconstruct I_s . However, the cycle-consistency brings more uncertainty of age mappings and cannot guarantee the bijective age mapping, especially when the gap between c_s and c_t becomes large.

To address this problem, we propose to model the age progression/regression as one invertible process using an invertible neural network [Dinh *et al.*, 2017]. Intuitively, the source age condition c_t could be extracted from the input face during performing the age progression/regression with the target age condition c_t . In this way, the flexibility of cGANs-based methods can be still enjoyed. Therefore, the invertible age

progression/regression process from the source age group s to target age group t can be formulated as follows:

$$\mathbf{I}'_t, \mathbf{c}'_s = G(\mathbf{I}_s, \mathbf{c}_t). \quad (2)$$

The inference phase of AgeFlow is illustrated in Figure 2. Given an input face \mathbf{I}_s , instead of directly modeling the patterns in the pixel space, the proposed method achieves age progression/regression on the latent vectors $\mathbf{z} = G(\mathbf{I})$ due to its reversibility, where G is a trained generative flow-based model [Kingma and Dhariwal, 2018]. Furthermore, we use a prior generator to produce the distribution of age condition as the conventional one-hot condition in the context of cGANs-based methods to keep consistency. This is because we found that the distribution contains more information than simple one-hot encoding that is too restrictive for the model to learn. In summary, the invertible conditional translation can be formulated as:

$$\mathbf{z}'_t, \boldsymbol{\mu}'_s, \boldsymbol{\sigma}'_s = T(\mathbf{z}_s, \boldsymbol{\mu}_t, \boldsymbol{\sigma}_t), \quad (3)$$

where T is the proposed invertible conditional translation module, and $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are the parameters of the Gaussian distribution from the prior generator for the age condition. Note that the decoder can easily reconstruct the learned latent vector back to the image domain through the inverse of the same encoder network, *i.e.*, G^{-1} . In other words, the cycle-consistency is achieved by the invertible network rather than reconstructing the input faces by another model.

2.2 Network Architecture

As shown in Figures 2 and 3, our AgeFlow consists of three key components to achieve an invertible conditional age progression/regression: 1) a flow-based encoder G that maps the input images into latent space and the decoder that is inverse function G^{-1} ; 2) an invertible conditional translation module T that achieves age progression/regression in the latent space; and 3) a discriminator to force T to generate age progressed/regressed latent vectors indistinguishable from real samples in the GAN framework [Goodfellow *et al.*, 2014] taking advantage of massive unpaired age data. Here we detail the network architectures of these three components.

Flow-based Generator. Flow-based generative models or GLOW [Kingma and Dhariwal, 2018] typically set up a bijective mapping between the image and latent space with invertible networks, which formulates the encoding process as $\mathbf{z} = G(\mathbf{I})$ and the generation one as the reverse procedure $\mathbf{I} = G^{-1}(\mathbf{z})$, where \mathbf{z} has a prior Gaussian distribution $p(\mathbf{z})$ and G is the invertible encoder. Due to the optimization of exact log-likelihood objective, flow-based models provide direct interpolations between data points and meaningful modifications of existing data points; in contrast, there is no encoder in GANs to model the full data distribution while VAEs maximize the evidence lower bound but generate blurry images. Here we adopt the similar structure of GLOW [Kingma and Dhariwal, 2018] for modeling the latent space with a standard Gaussian distribution. Specifically, as shown in Figure 2, there are $l \times k$ normalizing flows, each of which includes ActNorm, Additive Coupling, and invertible 1×1 convolutional layers, where the number of filters increases

twice and the size of feature maps decreases half for every k flows at each l scale. In this paper, we adopt 6×32 for input images with the size of 256×256 . Therefore, the loss function (maximum likelihood estimation) to optimize G can be written as:

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{z} \sim p_\theta(\mathbf{z})} \left[\log p_\theta(\mathbf{z}) + \log \left| \det \frac{dG}{d\mathbf{I}} \right| \right]. \quad (4)$$

Once trained, G is fixed to extract the latent vectors and the inverse of the trained G , G^{-1} , also serves as the decoder that maps the latent vector back to the image domain. In doing so, we can only train a small invertible translation network in the latent space to achieve age progression/regression, avoiding considerably computational cost of directly training invertible network from source age to target one in the image domain.

Invertible Conditional Translation Module. Given a latent vector \mathbf{z}_s from the source age group \mathbf{c}_s , the age progression/regression is to translate \mathbf{z}_s to \mathbf{z}_t at the target age \mathbf{c}_t . As illustrated in Figure 2, the proposed invertible conditional translation module (ICTM) T leverages the invertible network to achieve both age progression and regression, where the cycle consistency can be easily achieved due to its reversibility. In a sense, the inverse function of T plays a similar role in [Song *et al.*, 2018; Li *et al.*, 2020b] that achieves age regression independent of age progression. To further keep the flexibility of the cGANs-based method that achieves age progression/regression with only changing the target age condition, the translation module concatenates the latent vector \mathbf{z}_s with the distributions of target age condition estimated from the prior generator to achieve conditional translation. The reason why we adopt the distribution of age conditions is that the one-hot encoding used in previous works is too restrictive for the model to keep such cycle consistency and extract the age information from the input faces. Specifically, ICTM consists of m flows similar to previous invertible networks [Dinh *et al.*, 2017]. Each flow contains two convolutional layers, a channel attention module to enable the model to focus on changing the necessary latent variable, and a convolutional layer whose weights are initialized by zero to perform an identity function at the beginning of training. Finally, the progressed/regressed latent vectors and distributions of source age condition can be reconstructed from the two-way outputs of the ICTM, as indicated in Eq. (3).

Discriminator. We adopt a simple multi-layer perceptron (MLP) with two linear layers of 512 neurons as the discriminator D in our framework, each of which is followed by a spectral normalization [Miyato *et al.*, 2018] and a leaky ReLU activation whose negative slope is 0.2. At the bottom, we also append a linear layer with 1 neuron to output the confidence of GANs and another linear layer of n neurons for age classification to improve the age accuracy of our method, where n is the number of age groups.

2.3 Loss Function

The overall loss function to optimize the proposed ICTM contains four components for age progression/regression: 1) **Attribute-aware Knowledge Distillation Loss** aims to distill the manipulation direction of latent vectors for the

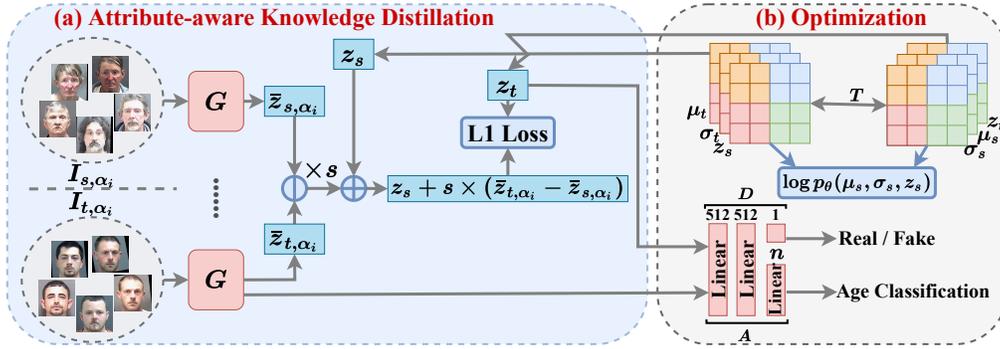


Figure 3: Training illustration of the ICTM. The average vectors z_{i, α_i} for each attribute pair are first computed from a set of images to form the manipulation direction, where i and α_i denotes the age group and facial attribute, respectively. They then become the label of ICTM to guide the directions of age progression/regression, *i.e.*, distilling the knowledge of the dataset to ICTM. There also includes GANs and age classification task for better image quality and age accuracy with the maximum log-likelihood estimation to keep the cycle consistency.

ICTM, which can preserve the flexibility of original cGANs-based methods with the facial attributes unchanged; 2) **Adversarial Loss** intends to force the ICTM to produce progressed/regressed latent variables indistinguishable from real ones; 3) **Age Classification Loss** expects to improve the aging accuracy, and 4) **Consistency Loss** seeks to regularize the ICTM with the cycle consistency.

Attribute-aware Knowledge Distillation Loss. The flow-based model can directly interpolate at the learned latent space and modify the data points to generate photorealistic images. For example, it can achieve image translation by applying the manipulation direction to the latent vector z_s of one given image, *i.e.* adding an average vector computed from a few images of the binary classes, which can be formulated as follows:

$$z'_t = z_s + s \times (z_{\text{pos}} - z_{\text{neg}}), \quad (5)$$

where s is a scale factor controlling the degree of manipulating z_s , and z_{pos} and z_{neg} are the average latent vectors computed from the sets of positive and negative samples, respectively. However, the original GLOW is not directly applicable to age progression/regression for the following reasons: 1) the facial attributes should be preserved during age progression/regression and 2) it cannot achieve conditional age progression/regression without the age label as prior and fails to ensure satisfied age accuracy; see the experiments for more evidence. To push the ICTM to manipulate the latent variables towards the direction of age progression/regression while keeping other attributes unchanged, we propose to use the learned manipulation direction to guide the ICTM in a knowledge distillation manner. Figure 3(a) shows the details of proposed attribute-aware knowledge distillation. Specifically, to preserve the facial attributes, we first compute the average vectors for each attribute pair from a set of images. By replacing the average latent vectors in Eq. (5), the consistency of facial attributes can be easily achieved. In addition, the learned direction can be further used as the supervision signals of ICTM optimized by the L1 distance. In summary, the knowledge distillation loss can be written as:

$$\mathcal{L}_{\text{akd}} = \mathbb{E}_{z_s} \left| z'_t - (z_s + s \times (\bar{z}_{t, \alpha_i} - \bar{z}_{s, \alpha_i})) \right|, \quad (6)$$

where s is a hyperparameter to achieve better age accuracy, and α_i represents one attribute of z_s .

Adversarial Loss. In this paper, we employ least squares GANs [Mao *et al.*, 2017] to optimize our discriminator in the latent space due to its stability, which can be defined as follows:

$$\mathcal{L}_{\text{al}} = \frac{1}{2} \mathbb{E}_{z_s} \left[(D(z'_t) - 1)^2 \right]. \quad (7)$$

Age Classification Loss. To improve the age accuracy of the generated faces, we also include an age classification loss in the final loss function, which can be defined as:

$$\mathcal{L}_{\text{acl}} = \mathbb{E}_{z_s} \left[\ell(A(z'_t), t) \right], \quad (8)$$

where ℓ is the cross-entropy loss and A denotes the age classifier as shown in Figure 3.

Consistency Loss. To keep the cycle consistency, the ICTM is supposed to output meaningful distribution of source age condition extracted from the input image. Therefore, we adopt a consistency loss that uses a standard Gaussian distribution to achieve the maximum log-likelihood estimation, which is formulated as:

$$\mathcal{L}_{\text{cl}} = -\mathbb{E}_{z_s \sim p_{\theta}(z_s)} [\log p_{\theta}(\mu'_s, \sigma'_s, z_s)]. \quad (9)$$

Due to the reversibility of invertible networks, we only need to maintain such cycle consistency for the output of ICTM. Besides, the input z_s can be easily reconstructed without any information loss.

Final Loss. In summary, the age progression/regression process requires age progressed/regressed latent variable to meet the following requirements: 1) indistinguishable from the real one and thus can be decoded to a photo-realistic face; 2) belongs to the target age group; and 3) has the same identity as the input one while keeping other facial attributes unchanged. Therefore, the final loss function to optimize ICTM is expressed as:

$$\mathcal{L}_T = \lambda_{\text{akd}} \mathcal{L}_{\text{akd}} + \lambda_{\text{al}} \mathcal{L}_{\text{al}} + \lambda_{\text{acl}} \mathcal{L}_{\text{acl}} + \lambda_{\text{cl}} \mathcal{L}_{\text{cl}}. \quad (10)$$

The loss function to optimize discriminator D is defined as:

$$\mathcal{L}_D = \frac{1}{2} \mathbb{E}_{z_t} (D(z_t) - 1)^2 + \frac{1}{2} \mathbb{E}_{z_s} D(z'_t)^2 + \lambda_{\text{acl}}^D \mathcal{L}_{\text{acl}}, \quad (11)$$

where z'_t is the learned latent vector belonging to target age group. During the training, the ICTM and discriminator D are updated alternatively.

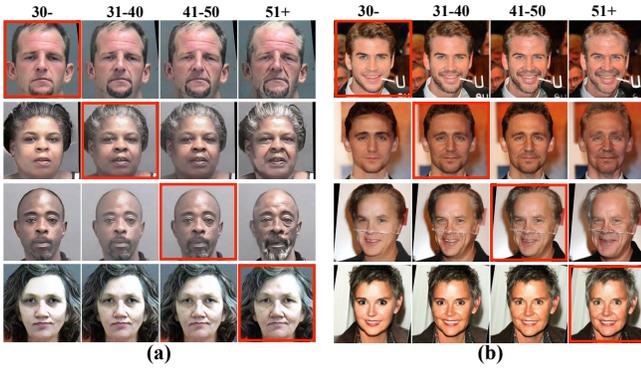


Figure 4: The generated faces by AgeFlow on (a) MORPH and (b) CACD. Red boxes indicate input faces for progression/regression.

3 Experiments

3.1 Implementation Details

Data Collection. We conducted experiments on two benchmark age datasets: **MORPH** [Ricanek and Tesafaye, 2006] and **CACD** [Chen *et al.*, 2015]. We also adopted **FG-NET** and **CelebA** [Liu *et al.*, 2015] as external testing sets to validate the model trained with CACD for a fair comparison. All face images were cropped and aligned based on landmarks detected by MTCNN [Zhang *et al.*, 2016] with an image size of 256×256 . We randomly divided the dataset into two parts without identities overlapping: 80% for training and the remaining for testing. Following the mainstream works [Yang *et al.*, 2018; Liu *et al.*, 2019; Li *et al.*, 2020b; Huang *et al.*, 2021a], we divided the face images into $n = 4$ age groups, *i.e.*, 30-, 31-40, 41-50, 50+.

Training Details. All models are implemented by PyTorch and trained by Adam optimizer with a fixed learning rate of 10^{-5} for the GLOW model and ICTM, and 10^{-4} for the discriminator. In addition, due to the limited GPU memory, we trained all models with a batch size of 16 on 4 NVIDIA V100 GPUs and the parameters are updated for every 4 iterations, equal to a batch size of 64. We first trained the GLOW model on CelebA [Liu *et al.*, 2015] thanks to its diversity and a huge amount of face images for 1M iterations and then finetuned the models on each dataset with only 50K iterations. ICTM contains $m = 32$ flows. The hyperparameters in the final loss are empirically set as follows: λ_{al} was 1; λ_{cl} was 0.01; λ_{acl} was 1; λ_{acl}^D was 0.1; and λ_{akl} was 1. The s in the knowledge distilling loss is set as 1.4 for MORPH and 1.8 for CACD. Here we emphasize that, to our best knowledge, we are the first work that has successfully reproduced GLOW from scratch with a high resolution of 256×256 . We also highlight that the pre-trained model will be very helpful for future works in the research community.

3.2 Qualitative Results

Figure 4 shows some generated faces by our AgeFlow from different ages to the other three age groups. Remarkably, the generated faces are photo-realistic and rendered with natural aging/rejuvenation effects; *e.g.*, the skins are tightened at young ages while the hair becomes white at old ages. Most

importantly, all faces present same identities as input faces, indicating that identity consistency has been well preserved.

Figure 5 shows the quantitative comparison between AgeFlow and other recent works including CAEE [Zhang *et al.*, 2017], IPCGAN [Wang *et al.*, 2018], PAG-GAN [Yang *et al.*, 2018], GLCA-GAN [Li *et al.*, 2018], Dual AcGAN [Li *et al.*, 2020b], and A3GAN [Liu *et al.*, 2019]. Markedly, our results have the highest image quality with natural effects and fewer artifacts. This is a direct result of the combination of flow-based models and GANs. Specifically, although the GANs-based method can learn good translation patterns, it suffers from unstable training and mode collapse, which results in strong artifacts and blurring in the generated images. However, the flow-based models optimize the exact log-likelihood objective function and thus are better at learning the underlying distribution of the dataset. In addition, thanks to the reversibility of flow-based models, the ghost artifacts would, as expected, not appear as we only perform GANs in the latent space instead of the image space. We highlight that the results of baseline methods are directly referred from their published papers to avoid any bias resulting from our self-implementation.

3.3 Quantitative Evaluations

We conduct qualitative comparisons with prior work including CAEE [Zhang *et al.*, 2017], IPCGAN [Wang *et al.*, 2018], and WGLCA-GAN [Li *et al.*, 2019] with an ablation study. We employ two widely-used quantitative metrics to evaluate age progression/regression methods: age accuracy and identity preservation, *i.e.*, the progressed/regressed faces should ideally belong to the target age group with identities preserved. We also evaluated our method in preserving the facial attributes of input faces to demonstrate the effectiveness of the proposed attribute-aware knowledge distillation loss.

Age Accuracy. Following [He *et al.*, 2019], we trained a ResNet-101 for each dataset as the age classifier to measure the age accuracy for different methods. Specifically, we translate the faces into the other three age groups, and we used the percentage of the faces falling into the exact target age group as the age accuracy, where the higher is the better. Table 1 presents the average age accuracy by different methods in each group. Our methods outperform other baseline methods by a large margin among all four age groups, demonstrating its capability of aging/rejuvenating faces with strong and natural aging effects.

Identity Preservation. We performed the face verification task to evaluate the identity preservation of different methods. Specifically, we used the pre-trained ResNet-101 face recognition model from ArcFace [Deng *et al.*, 2019] to calculate the cosine similarity between the input and generated images. We reported the average cosine similarity of each pair of images as identity preservation. Here we emphasize that the verification rate is essentially computed by setting a threshold for the cosine similarity; that is, there is no difference between cosine similarity and face verification rate. The results are shown in Table 1. Remarkably, our methods still achieve the highest scores at four age groups, confirming that our methods can achieve a satisfying age accuracy with the

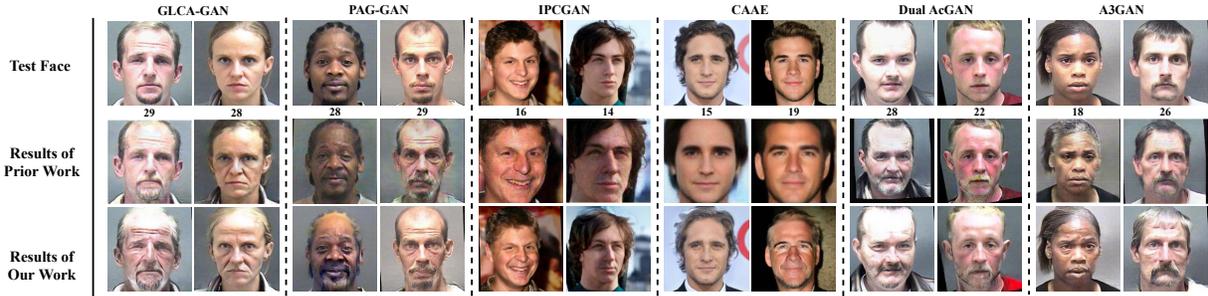


Figure 5: Performance comparison with prior work on the MORPH and CACD datasets. The three rows are the input young faces with their real age labels underneath, the results of prior work, and our results in the same age groups (51+), respectively.

MORPH					CACD				
Age group	30 -	31 - 40	41 - 50	50 +	Age group	30 -	31 - 40	41 - 50	50 +
Age Accuracy (%)					Age Accuracy (%)				
CAAE	56.71	56.17	45.58	48.48	CAAE	57.50	42.60	50.51	24.97
IPCGAN	75.71	80.30	48.91	53.71	IPCGAN	71.84	52.99	60.07	77.61
WGLCA-GAN	80.19	77.10	81.12	85.38	WGLCA-GAN	78.05	72.48	62.54	81.48
AgeFlow (ours)	87.19	91.85	90.68	90.71	AgeFlow (ours)	86.49	78.06	85.91	88.15
Cosine Similarity					Cosine Similarity				
CAAE	0.265	0.289	0.279	0.220	CAAE	0.092	0.101	0.091	0.075
IPCGAN	0.567	0.609	0.610	0.467	IPCGAN	0.445	0.554	0.569	0.462
WGLCA-GAN	0.671	0.792	0.764	0.614	WGLCA-GAN	0.752	0.803	0.769	0.758
AgeFlow (ours)	0.897	0.923	0.903	0.767	AgeFlow (ours)	0.905	0.926	0.919	0.847

Table 1: Quantitative results of different methods on MORPH and CACD datasets. The best results are shown in bold.

Age Group	Gender						Race		
	MORPH			CACD			MORPH		
	31 - 40	41 - 50	51 +	31 - 40	41 - 50	51 +	31 - 40	41 - 50	51 +
CAAE	94.53	94.59	93.37	92.96	89.11	83.78	97.01	96.73	93.84
IPCGAN	97.30	97.48	95.14	94.96	97.42	95.49	96.87	97.10	97.05
WGLCA-GAN	97.71	98.95	96.20	98.14	95.64	95.69	98.53	98.26	96.28
AgeFlow (ours)	98.19	99.26	98.21	98.63	98.02	98.49	98.85	98.80	97.72

Table 2: Attribute preservation for different methods on MORPH and CACD datasets. The best results are shown in bold.

identities consistently preserved.

Attribute Preservation. Similar to age accuracy, we trained a classifier for each facial attribute of each dataset and reported the aging process starting from 30- as it is the hardest case in age progression/regression. Specifically, we evaluate the gender and race attributes for MORPH and only gender for CACD as there are almost all white people in CACD dataset. We argue that gender and race are the two most important facial attributes affecting face age synthesis. Other attributes such as hairstyle and expression are rarely discussed and not available in the benchmarked datasets. Table 2 presents the attribute preservation for different methods. Clearly, our method outperforms the baseline methods on the two datasets for different facial attributes. It is benefitting from the merit of invertible neural networks—bijective mapping to preserve age-irrelevant information, which could best reduce the uncertainty of face transformations between two

age groups. We note that the proposed method only needs the attribute labels during training compared to [Liu *et al.*, 2019] that injects the attributes into the generator, therefore, it enjoys the flexibility of cGANs-based methods while consistently preserving facial attributes.

3.4 Ablation Study

To further validate the effectiveness of AgeFlow, we performed an ablation study for the following variants: 1) GLOW: the trained GLOW model is used to manipulate the faces towards the target age group; 2) GLOW+attribute: we leverage the facial attributes to demonstrate that the prior attribute information can improve the attribute consistency during age progression/regression; 3) AgeFlow w/o ICTM: we replace ICTM with a series of convolutional layers like cGANs to achieve the transformation in the latent space; and 4) AgeFlow: our proposed method distills the manipulation

MORPH					CACD				
Age group	30 -	31 - 40	41 - 50	50 +	Age group	30 -	31 - 40	41 - 50	50 +
Age Accuracy (%)					Age Accuracy (%)				
GLOW	47.92	52.75	60.46	65.23	GLOW	49.92	49.64	52.61	32.13
+Attribute	61.86	54.51	64.17	35.67	+Attribute	67.35	42.39	55.39	46.95
AgeFlow	87.19	91.85	90.68	90.71	AgeFlow	86.49	78.06	85.91	88.15
w/o ICTM	80.12	87.27	84.92	88.37	w/o ICTM	78.05	67.51	80.23	79.98
Cosine Similarity					Cosine Similarity				
GLOW	0.809	0.772	0.912	0.721	GLOW	0.811	0.865	0.871	0.773
+Attribute	0.853	0.917	0.915	0.752	+Attribute	0.895	0.910	0.925	0.796
AgeFlow	0.897	0.923	0.903	0.767	AgeFlow	0.905	0.926	0.919	0.847
w/o ICTM	0.870	0.889	0.873	0.739	w/o ICTM	0.877	0.887	0.901	0.749

Table 3: Quantitative results for different variants of AgeFlow on MORPH and CACD datasets. The best results are shown in bold.

Age Group	Gender						Race		
	MORPH			CACD			MORPH		
	31 - 40	41 - 50	51 +	31 - 40	41 - 50	51 +	31 - 40	41 - 50	51 +
GLOW	95.02	94.89	94.60	94.38	95.77	94.36	97.34	96.88	96.45
+Attribute	97.62	98.95	98.22	98.68	98.01	98.63	98.72	98.83	97.09
AgeFlow	98.19	99.26	98.21	98.63	98.02	98.49	98.85	98.80	97.72
w/o ICTM	96.15	98.48	96.20	98.14	96.66	98.30	99.41	97.71	96.75

Table 4: Attribute preservation for different variants of AgeFlow on MORPH and CACD dataset. The best results are shown in bold.

direction learned from the dataset to the ICTM.

Table 3 shows the age accuracy and identity preservation while Table 4 presents the results of facial attribute preservation during age progression/regression. It can be observed that GLOW and its variant GLOW+attributes failed to achieve a satisfying age accuracy. This is because GLOW achieves age progression/regression by interpolating on the latent space, and they control the degree of age progression/regression by applying the manipulation direction. In a sense, different faces share the same patterns, which thus causes a low age accuracy. However, the attributes can be preserved by GLOW+attribute as it leverages the prior attributes information. Without ICTM, it fails to preserve the identity with the attributes changed, since the non-invertible introduces strong uncertainty to the age progression/regression process. Figure 6 shows that the attributes were changed by the original GLOW but preserved by other variants. However, their aging/rejuvenation effects are not so strong as the proposed AgeFlow, since GLOW and GLOW+attribute are limited to the interpolations of latent space, and w/o ICTM cannot ensure an invertible between input and generated faces. The results of the ablation study confirm the effectiveness of the prior attribution information and the proposed ICTM.

4 Conclusion

In this paper, we proposed a novel framework, called AgeFlow, to integrate the advantages of both flow-based models and GANs, which enables age progression and regression to be invertible and achieves face transformation in the latent space with GANs and attribute-aware knowledge distillation

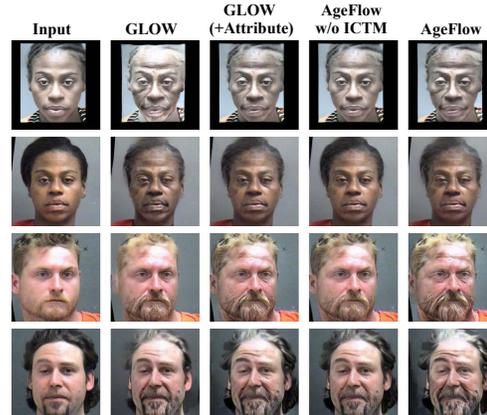


Figure 6: Qualitative comparisons of different variants of AgeFlow.

loss. Extensive experiments demonstrate the superior performance over recently proposed baseline methods. Our work could shed new light on training invertible neural networks with high-resolution images and using GANs in the latent space learned from flow-based methods.

Acknowledgements

This work was supported in part by Shanghai Sailing Program (No. 21YF1402800), National Key Research and Development Program of China (No. 2018YFB1305104), Shanghai Municipal of Science and Technology Project (No. 20JC1419500), Shanghai Municipal Science and Technology Major Project (No. 2018SHZDZX01), and ZJ Lab.

References

- [Chen *et al.*, 2015] Bor-Chun Chen, Chu-Song Chen, and Winston H Hsu. Face recognition and retrieval using cross-age reference coding with cross-age celebrity dataset. *IEEE TMM*, 2015.
- [Deng *et al.*, 2019] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- [Dinh *et al.*, 2015] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: Non-linear independent components estimation. In *ICLR*, 2015.
- [Dinh *et al.*, 2017] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *ICLR*, 2017.
- [Goodfellow *et al.*, 2014] 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.
- [He *et al.*, 2019] Zhenliang He, Meina Kan, Shiguang Shan, and Xilin Chen. S2GAN: Share aging factors across ages and share aging trends among individuals. In *ICCV*, 2019.
- [Huang *et al.*, 2021a] Zhizhong Huang, Shouzheng Chen, Junping Zhang, and Hongming Shan. PFA-GAN: Progressive face aging with generative adversarial network. *IEEE TIFS*, 2021.
- [Huang *et al.*, 2021b] Zhizhong Huang, Junping Zhang, and Hongming Shan. Routinggan: Routing age progression and regression with disentangled learning. In *ICASSP*, 2021.
- [Huang *et al.*, 2021c] Zhizhong Huang, Junping Zhang, and Hongming Shan. When age-invariant face recognition meets face age synthesis: A multi-task learning framework. In *CVPR*, 2021.
- [Kingma and Dhariwal, 2018] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *NeurIPS*, 2018.
- [Li *et al.*, 2018] Peipei Li, Yibo Hu, Qi Li, Ran He, and Zhenan Sun. Global and local consistent age generative adversarial networks. In *ICPR*, 2018.
- [Li *et al.*, 2019] Peipei Li, Yibo Hu, Ran He, and Zhenan Sun. Global and local consistent wavelet-domain age synthesis. *IEEE TIFS*, 2019.
- [Li *et al.*, 2020a] Peipei Li, Huaibo Huang, Yibo Hu, Xiang Wu, Ran He, and Zhenan Sun. Hierarchical face aging through disentangled latent characteristics. In *ECCV*, 2020.
- [Li *et al.*, 2020b] Qi Li, Yunfan Liu, and Zhenan Sun. Age progression and regression with spatial attention modules. In *AAAI*, 2020.
- [Liu *et al.*, 2015] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015.
- [Liu *et al.*, 2019] Yunfan Liu, Qi Li, and Zhenan Sun. Attribute-aware face aging with wavelet-based generative adversarial networks. In *CVPR*, 2019.
- [Mao *et al.*, 2017] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017.
- [Mirza and Osindero, 2014] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [Miyato *et al.*, 2018] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.
- [Or-El *et al.*, 2020] Roy Or-El, Soumyadip Sengupta, Ohad Fried, Eli Shechtman, and Ira Kemelmacher-Shlizerman. Lifespan age transformation synthesis. In *ECCV*, 2020.
- [Park *et al.*, 2010] Unsang Park, Yiyang Tong, and Anil K Jain. Age-invariant face recognition. *IEEE TPAMI*, 2010.
- [Ricanek and Tesafaye, 2006] Karl Ricanek and Tamirat Tesafaye. MORPH: A longitudinal image database of normal adult age-progression. In *FG*, 2006.
- [Song *et al.*, 2018] Jingkuan Song, Jingqiu Zhang, Lianli Gao, Xianglong Liu, and Heng Tao Shen. Dual conditional gans for face aging and rejuvenation. In *IJCAI*, 2018.
- [Wang *et al.*, 2016] Wei Wang, Zhen Cui, Yan Yan, Jiashi Feng, Shuicheng Yan, Xiangbo Shu, and Nicu Sebe. Recurrent face aging. In *CVPR*, 2016.
- [Wang *et al.*, 2018] Zongwei Wang, Xu Tang, Weixin Luo, and Shenghua Gao. Face aging with identity-preserved conditional generative adversarial networks. In *CVPR*, 2018.
- [Yang *et al.*, 2018] Hongyu Yang, Di Huang, Yunhong Wang, and Anil K Jain. Learning face age progression: A pyramid architecture of GANs. In *CVPR*, 2018.
- [Zhang *et al.*, 2016] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE SPL*, 2016.
- [Zhang *et al.*, 2017] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *CVPR*, 2017.
- [Zhu *et al.*, 2017] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.