

Pluralistic Aging Diffusion Autoencoder

Peipei Li¹, Rui Wang¹, Huaibo Huang², Ran He², Zhaofeng He^{1*}

¹Beijing University of Posts and Telecommunications

²CRIPAC&MAIS, Institute of Automation, Chinese Academy of Sciences

{lpeipei, wr_bupt, zhaofenghe}@bupt.edu.cn, huaibo.huang@cripac.ia.ac.cn, rhe@nlpr.ia.ac.cn

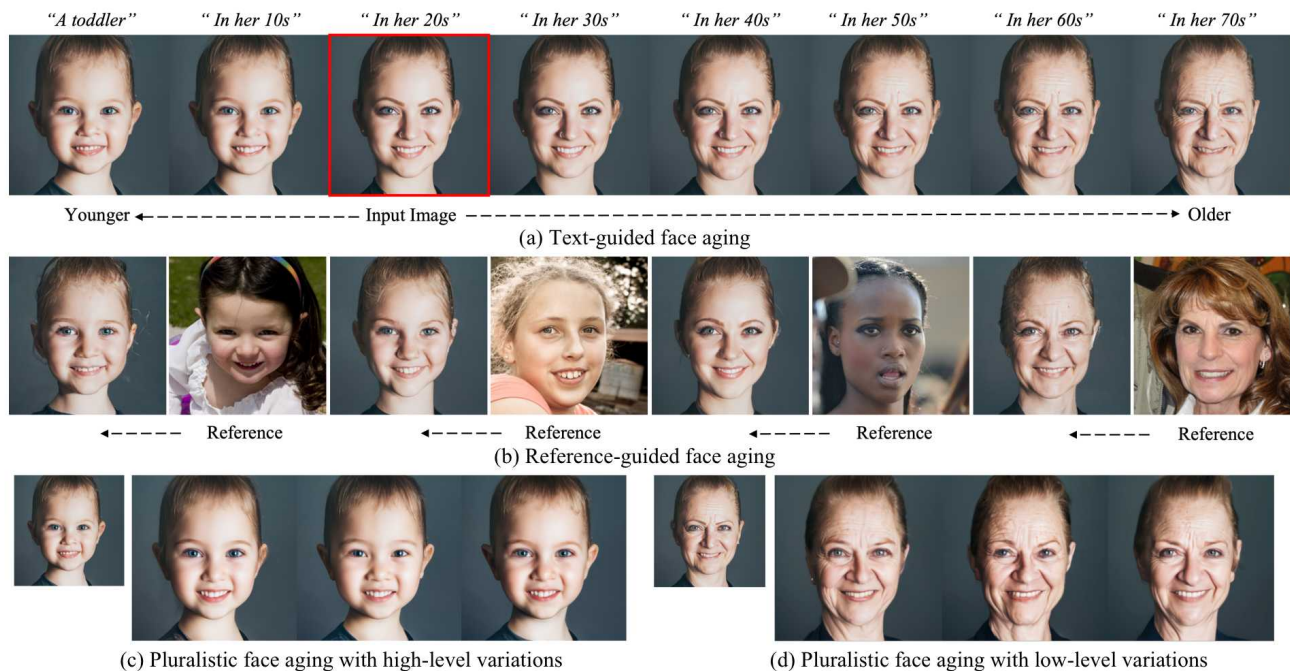


Figure 1: Our single framework enables various pluralistic face aging tasks, including (a) Text-guided face aging (the first row), (b) Reference-guided face aging (the second row), (c) Pluralistic face aging with high-level variations (the bottom) and (d) Pluralistic face aging with low-level variations (the bottom).

Abstract

Face aging is an ill-posed problem because multiple plausible aging patterns may correspond to a given input. Most existing methods often produce one deterministic estimation. This paper proposes a novel CLIP-driven Pluralistic Aging Diffusion Autoencoder (PADA) to enhance the diversity of aging patterns. First, we employ diffusion models to generate diverse low-level aging details via a sequential denoising reverse process. Second, we present Probabilistic Aging Embedding (PAE) to capture diverse high-level aging patterns, which represents age information as probabilistic distributions in the common CLIP latent space. A text-

guided KL-divergence loss is designed to guide this learning. Our method can achieve pluralistic face aging conditioned on open-world aging texts and arbitrary unseen face images. Qualitative and quantitative experiments demonstrate that our method can generate more diverse and high-quality plausible aging results.

1. Introduction

Face aging aims to model facial appearance changes across different ages meanwhile maintaining identity information. It is an ill-posed learning problem due to multiple plausible aging results for a given input. Given various aging images in the last row of Fig. 1, which one meets your

*Corresponding author

imagination of aging? Since human aging process is influenced by a variety of factors, including genetics and social environment, there may be significant differences in both general aging trends and local details. Pluralistic face aging aims to generate multiple and diverse plausible face aging images from a single input.

Deep generative models, such as generative adversarial networks (GANs) [9] and variational autoencoders (VAEs) [17], have shown impressive performance in terms of face aging [24, 28, 26, 1, 10]. Unfortunately, most previous methods can only produce one “optimal” aging pattern, which is inconsistent with human cognition. Recently, diffusion models [6, 16] show comparable or even better generation quality compared to GANs, which learn the reverse of a particular Markov diffusion process and cover the modes of data distribution better. Inspired by this, we intend to employ diffusion models to generate aging faces with low-level subtle stochastic variations, such as diverse wrinkles, as shown in the lower right of Fig. 1 (d).

In addition to low-level stochastic details, the aging process is accompanied by high-level age semantic changes, such as, getting fatter or thinner, getting darker or whiter, as shown in Fig. 1 (c). Previous face aging methods [28, 1, 10, 8] directly represent the target age as a deterministic point or direction in the latent space, ignoring the personalized age characteristics. So here comes a key challenge for pluralistic face aging: how to learn high-level age representations with stochastic variations. To address it, we draw support from the pre-trained CLIP [32] model and propose Probabilistic Aging Embedding (PAE), which represents age information as a distribution rather than a deterministic point. The intuition to leverage CLIP is illustrated in Fig. 2. In the well-aligned image-text latent space, there are likely to be multiple image-based age features for a coarse text-based age feature of “*Man’s face in his forties*”. Inspired by it, we attempt to model PAE in CLIP latent space to capture the stochastic high-level age semantics.

In this paper, we explore pluralistic face aging based on both text and image conditions. We propose a CLIP-driven Pluralistic Aging Diffusion Autoencoder (PADA) to simultaneously model low-level stochastic variations and high-level age semantic variations in the aging process. For the low-level age variations, our method is based on the diffusion model, which can generate stochastic low-level face details via a sequential denoising procedure. For the high-level age variations, we propose Probabilistic Aging Embedding (PAE) by representing the age information as probabilistic distributions in the CLIP space. Specifically, we represent the age information as a multivariate Gaussian distribution rather than a deterministic point, where the mean of the distribution indicates average age information, while the variance indicates the personalized aging patterns of the image. Then we feed back the PAE into the diffu-

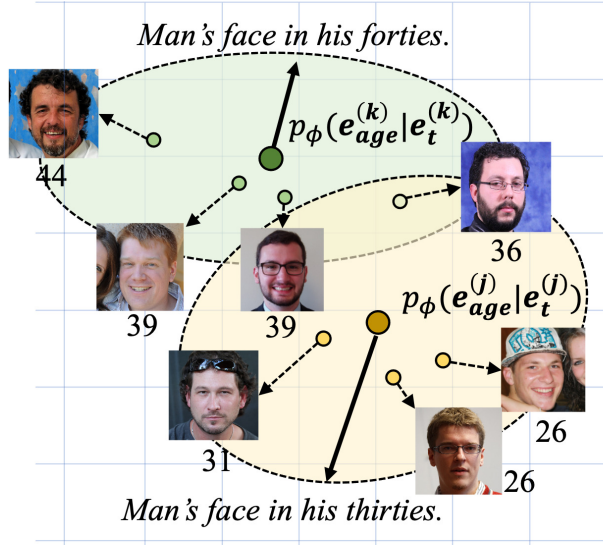


Figure 2: One-to-many correspondences between coarse text-based age feature and image-based age features in the CLIP latent space. The solid line indicates the text-based age feature while the dashed line indicates the image-based age features.

sion model with an adaptive modulation for pluralistic face aging. Since our goal is to learn the diverse aging patterns and achieve face aging with preservation of age-irrelevant information (i.e., identity and background), three types of losses are employed: 1) Text-guided KL-divergence loss; 2) Age fidelity loss; 3) Preservation loss. To summarize, our contributions are four-fold:

- We propose a novel CLIP-driven Pluralistic Aging Diffusion Autoencoder (PADA) for pluralistic face aging, which can generate diverse aging results with both high-level age semantic variations and low-level stochastic variations.
- Probabilistic Aging Embedding (PAE) is proposed in the common CLIP space to represent the diverse high-level aging patterns as probabilistic distribution, where a text-guided KL-divergence loss is employed to guide this learning.
- A more user-friendly interaction way for face aging is provided, which can achieve age manipulation conditioned on both the open-world age descriptions and arbitrary unseen face images in the wild.
- Extensive qualitative and quantitative experiments show that our method outperforms the state-of-the-art

aging methods and can generate plausible and diverse aging patterns.

2. Related Work

Face Aging is one of the most challenging components in modern face manipulation, which is an ill-posed problem with multiple aging results corresponding to the same input. Generative Adversarial Network (GAN) [9] has been successfully applied to face aging and made impressive results [8, 21, 22, 24, 26, 25, 23, 28, 1, 10]. LATS [28] successfully models both the texture transformation and shape deformation. SAM [1] is based on a pre-trained StyleGAN [12] and can generate high-quality aging results. RA-GAN [26] proposes a personalized self-guidance module, which leverages the interactions between identity and target age to learn the personalized age features. CUSP [8] disentangles the style and content of the input, providing structure modifications with relevant details unchanged. Variational Autoencoder (VAE) shows promising ability in generating results with interpretability [17] and is employed for face aging [38, 24]. Li et al. [24] propose a disentangled adversarial autoencoder (DAAE) for face aging, which disentangles the images into three independent factors, including age, identity and extraneous information. All in all, although there may be many reasonable possibilities, the most existing methods produce only one “optimal” estimation for each input. Inspired by probabilistic embedding learning [20, 4, 3], we intend to utilize pre-trained CLIP to model Probabilistic Aging Embedding (PAE) when paired aging data is unavailable.

Recently, **Diffusion Probabilistic Models (DPM)** have achieved state-of-the-art results in generation quality as well as in sample quality [6, 16, 27, 14], which consist of a forward (or inference) Markovian diffusion process q and a learned reverse (or generative) diffusion process $p_{\theta}:q(x_{1:T}|x_0) = \prod_{t=1}^T q(x_t|x_{t-1}), p_{\theta}(x_{0:T}) = p(x_T) \prod_{t=1}^T p_{\theta}(x_{t-1}|x_t)$. More formally, we define the forward diffusion process at time t by gradually adding Gaussian noise to the input x_0 as $q(x_t|x_{t-1}) = \mathcal{N}(x_t|\sqrt{\alpha_t}x_{t-1}, (1-\alpha_t)I)$; while the inverse process by gradually removing the noise from the Gaussian noise $p_{\theta}(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}|\mu_{\theta}(x_t, t), \sigma_t^2 I)$. Thus, DDPMs require simulating a Markov chain with many iterations to produce a high quality sample, which limits its generation efficiency. Denoising Diffusion Implicit Models (DDIMs) [35] are then proposed to accelerate sampling, which share the same training objective function with DDPMs, but employ non-Markovian diffusion processes. The deterministic generative process and the following inference distribution are shown as follows:

$$x_{t-1} = \sqrt{\alpha_{t-1}}f_{\theta}(x_t, t) + \sqrt{1-\alpha_{t-1}}\epsilon_{\theta}^t(x_t), \quad (1)$$

$$q(x_{t-1}|x_t, x_0) = \mathcal{N}\left(\sqrt{\alpha_{t-1}}x_0 + \sqrt{1-\alpha_{t-1}}\frac{x_t - \sqrt{\alpha_t}x_0}{\sqrt{1-\alpha_t}}, \mathbf{0}\right), \quad (2)$$

where $f_{\theta}(x_t, t)$ is the prediction of x_0 at timestep t given predicted noise $\epsilon_{\theta}^t(x_t)$:

$$f_{\theta}(x_t, t) = \frac{x_t - \sqrt{1-\alpha_t}\epsilon_{\theta}^t(x_t)}{\sqrt{\alpha_t}}. \quad (3)$$

To learn a meaningful latent space, DiffAE [31] trains DDPM with an extra encoder, which embeds the input into a latent vector to guide the reverse diffusion process.

With the development of the powerful cross-modal visual and language model CLIP [32], many recent efforts start to study **CLIP-Guided Image Generation** [30, 18, 36, 37, 7, 14, 2]. CLIPStyler [18] utilizes the well-aligned CLIP latent space for high-quality style transfer. HairCLIP [37] is based on pre-trained StyleGAN [13] and CLIP for a more user-friendly design of hairstyles. However, there is no existing method fully leverage the CLIP latent space for lifespan aging. In this paper, we explore CLIP latent space to provide diverse age representations for face aging.

3. Proposed Method

3.1. Overview

We propose CLIP-driven Pluralistic Aging Diffusion Autoencoder (PADA), a conditional DDIM to generate multiple face aging results. Specifically, we aim to transform the source image x_0^{src} to a set of target aging results $\{(x_0^{tar})^{(i)}\}_{i=1}^N$ conditioned on either reference image x_0^{ref} or text description t^{ref} . x_0^{tar} is the target aging result with stochastic age variations, which is generated via an aging reverse process of DDIM:

$$p(x_{0:T}^{tar}|z^{src}, z^{age}) = p(x_T) \prod_{t=1}^T p(x_{t-1}^{tar}|x_t^{tar}, z^{src}, z^{age}). \quad (4)$$

Specifically, our aging reverse process is based on three parts, including a pre-trained conditional DDIM decoder $p(x_{t-1}^{tar}|x_t^{tar}, z^{src}, z^{age})$, a pre-trained semantic encoder $z^{src} = E_{sem}(x_0^{src})$, and a CLIP-guided age encoder $z^{age} = E_{age}(x_0^{ref}, t^{ref})$, where z^{src} is the semantic information of input image x_0^{src} ; and z^{age} is the stochastic age condition learned from reference image x_0^{ref} or text t^{ref} . An overview of our architecture is shown in Fig. 3.

Unlike previous methods [28, 1, 10, 8] that learn a determined age variable, we propose CLIP-guided age encoder E_{age} to learn the stochastic age condition z^{age} for pluralistic face aging. Concretely, leveraging the well-aligned text-image latent space of the pre-trained CLIP model, we extract the **Probabilistic Aging Embedding (PAE)**. Then, an adaptive modulation mechanism is utilized to translate

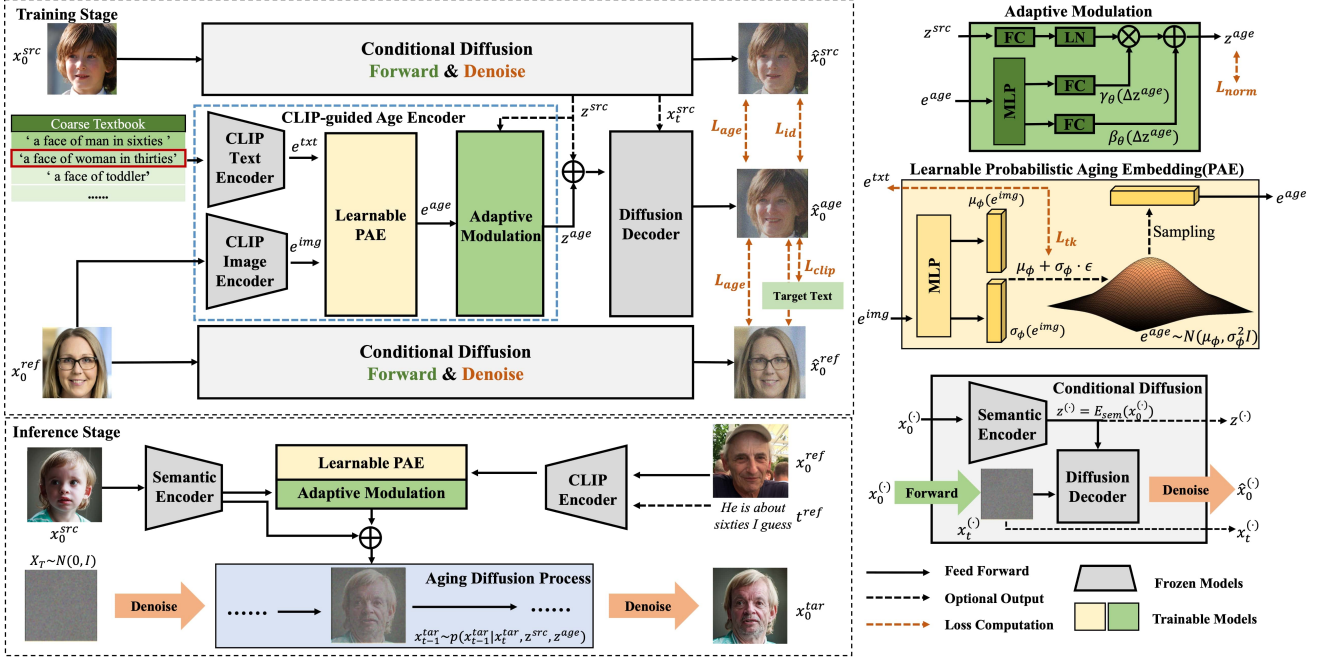


Figure 3: Overview of the proposed PADA, which consists of a conditional DDIM decoder, a semantic encoder, and a CLIP-guided age encoder. The semantic information z^{src} is extracted from the input image x_0^{src} via the semantic encoder. Meanwhile, the Probabilistic Aging Embedding (PAE) e^{age} is obtained by the CLIP-guided age encoder and translated to z^{age} in an adaptive manner. Based on z^{src} and z^{age} , we can generate pluralistic aging results via the conditional DDIM decoder. In the inference stage, e^{age} can also be sampled from text-based age prior.

PAE into the stochastic age condition z^{age} . Conditioned on both z^{age} and z^{src} , we can generate pluralistic aging results via the pre-trained conditional DDIM decoder. The training and inference algorithms are detailed in the supplementary materials.

3.2. Stochastic Age Condition

In this section, we introduce two key ingredients for the learning of stochastic age condition z^{age} : Probabilistic Aging Embedding (PAE) and Adaptive Modulation.

3.2.1 Probabilistic Aging Embedding

As illustrated in Fig. 2, in the well-aligned image-text latent space of the pre-trained CLIP [32] model, the cosine similarity of related images and texts is maximized, while that of unrelated images and texts is minimized. Since coarse text-based age features contain the average aging information (e.g., *Man's face in his forties*), rich personalized age-related fine features are assumed to be most likely distributed around it. Therefore, to obtain richer personalized age-related features, we propose Probabilistic Aging Embedding (PAE) based on well-aligned CLIP latent space.

Concretely, we represent our PAE as a multivariate

Gaussian distribution $\mathcal{N}(e^{age}; \mu_\phi, \sigma_\phi^2 I)$ in the CLIP latent space. The mean μ_ϕ indicates average age information, while the variance σ_ϕ indicates the personalized aging patterns from the reference image. We assume the posterior approximation $p_\phi(e^{age}|e^{img})$ follows an isotropic multivariate Gaussian:

$$p_\phi(e^{age}|e^{img}) = \mathcal{N}(e^{age}; \mu_\phi(e^{img}), \sigma_\phi^2(e^{img})I), \quad (5)$$

where e^{img} is extracted from the reference image x_0^{ref} via the CLIP image encoder. $\mu_\phi(\cdot)$ and $\sigma_\phi(\cdot)$ are implemented with a shared MLP-backbone and two head branches. Moreover, we introduce an age prior for PAE: $p(e^{age}) = \mathcal{N}(e^{txt}, I)$, where e^{txt} is extracted from the coarse age description t^{ref} via the CLIP text encoder.

Correspondingly, we use a text-guided KL-divergence loss for prior matching. The learned probabilistic aging embedding e^{age} is sampled via a reparameterization trick: $e^{age} = \mu_\phi + \sigma_\phi \odot \eta$, where $\eta \sim \mathcal{N}(0, I)$ and \odot refers to element-wise multiplication. To get richer aging features, we also directly model PAE in a text-driven manner: $e^{age} = e^{txt} + \epsilon \cdot \eta$, where ϵ is a hyperparameter for sampling intensity. We can prove that by careful selection of ϵ , probabilistic aging embeddings e^{age} can be reliably sampled on the hyper-sphere of CLIP latent space. The formal proof is

provided in the supplementary materials.

In this way, we can sample probabilistic aging embedding e^{age} from our learned distribution.

3.2.2 Adaptive Modulation

Our Adaptive Modulation is designed to translate the learned e^{age} into stochastic age condition z^{age} , which is composed of an MLP-based backbone and a modulation module. Concretely, for each of the sampled e^{age} , we first translate it into semantic latent space of the conditional DDIM: $\Delta z^{age} = MLP(e^{age})$. Then the aging information Δz^{age} is transferred into the stochastic age condition z^{age} by our modulation module, which can be formulated as:

$$z^{age} = \gamma_{\theta}(\Delta z^{age}) \frac{z^{src} - \mu^{src}}{\sigma^{src}} + \beta_{\theta}(\Delta z^{age}), \quad (6)$$

where μ^{src} and σ^{src} are the channel-wise mean and standard deviation of z^{src} , respectively. Here, we construct $\gamma_{\theta}(\cdot)$ and $\beta_{\theta}(\cdot)$ with 2 fully-connected layers. Conditioned on both z^{src} and z^{age} , we can generate pluralistic face aging results x_0^{tar} via the pre-trained DDIM decoder $p(x_{t-1}^{tar} | x_t^{tar}, z^{src}, z^{age})$.

3.3. Loss Functions

Our goal is to achieve pluralistic face aging with age-irrelevant information well-preserved. Therefore, three types of loss objectives are designed, including text-guided KL-divergence loss, age fidelity loss, and preservation loss. To avoid confusion, here we denote $\hat{x}_0^{(\cdot)}$ as the approximate reconstruction of $x_0^{(\cdot)}$ at timestep t following Eq.(3). For example, \hat{x}_0^{tar} means the approximate reconstruction of aging result x_0^{tar} at time step t .

Text-guided KL-divergence Loss. To prevent the learned variances from collapsing to zero, we explicitly constrain PAE to be close to a Gaussian distribution by introducing a text-guided KL-divergence loss L_{tKL} :

$$L_{tKL} = D_{KL}(\mathcal{N}(e^{age}; \mu_{\phi}, \sigma_{\phi}^2 I) || \mathcal{N}(e^{txt}, I)) \\ = \frac{1}{2} \sum_j^D \left((\mu_{\phi}^j - (e^{txt})^j)^2 - 1 + (\sigma_{\phi}^j)^2 - \log(\sigma_{\phi}^j)^2 \right), \quad (7)$$

where D denotes the channel of e^{age} . For training stability, we replace the Euclidean distance term with a cosine similarity term $-\cos(\mu_{\phi}, e^{txt})$ and a norm term $\|\mu_{\phi} - e^{img}\|_2$, which is optimally equivalent when considering CLIP spaces as hyperspheres. More details can be found in supplementary materials.

In practice, we normalize the features in CLIP latent space by their L_2 norm.

Age Fidelity Loss. To better maintain age fidelity, we employ an age prediction loss and a clip directional loss. Concretely, we first use a pre-trained age estimator DEX [34] to estimate the age of intermediate generation \hat{x}_0^{tar} in Eq.(3). Due to the blurry of \hat{x}_0^{tar} , we use an aging triplet loss for better age fidelity. For conciseness, we denote $f^{(\cdot)}$ as aging representation in the last fully-connected layer of the age predictor:

$$L_{age} = \max \{ \langle f^{src}, f^{tar} \rangle - \langle f^{ref}, f^{tar} \rangle + m, 0 \}, \quad (8)$$

where $\langle \cdot, \cdot \rangle$ refers to the cosine similarity and m refers to the margin.

To learn more aging details, we additionally employ a CLIP directional loss L_{clip} [30].

Preservation Loss. To better preserve the identity and age-irrelevant information, we employ identity loss L_{id} , norm loss L_{norm} , and reconstruction loss L_{rec} . Concretely, the identity preservation loss L_{id} is used to preserve the identity during generation, which is formulated as:

$$L_{id} = -\cos(R(\hat{x}_0^{src}), R(\hat{x}_0^{tar})), \quad (9)$$

where $R(\cdot)$ is output of the final fully-connected layer of the pre-trained ArcFace [5]. Norm loss L_{norm} is the regularization term to ensure the generation quality:

$$L_{norm} = \|z^{age}\|_2^2. \quad (10)$$

To ensure the age-irrelevant information unchanged, we introduce the reconstruction loss L_{rec} , which is defined as:

$$L_{rec} = \|\hat{x}_0^{src} - \hat{x}_0^{tar}\|_2^2. \quad (11)$$

The overall loss of PADA is formulated as:

$$L = L_{age} + \lambda_1 L_{clip} + \lambda_2 L_{tKL} + \lambda_3 L_{id} \\ + \lambda_4 L_{norm} + \lambda_5 L_{rec}, \quad (12)$$

where λ_i is hyperparameter controlling the weight of each loss. The details are described in Section 4.

4. Experiments

Datasets. We relabeled a new facial aging dataset based on the images of FFHQ [12], FFHQ-AT. We roughly divided the FFHQ into 7 age groups: 0-5, 6-15, 16-25, 26-35, 36-50, 51-70, 70+ and accordingly pre-defined 14 age-related descriptions, including ‘a toddler girl’s/boy’s face’, ‘woman’s/man’s face in her/his tens’, ‘woman’s/man’s face in her/his twenties’, ‘woman’s/man’s face in her/his thirties’, ‘woman’s/man’s face in her/his forties’, ‘woman’s/man’s face in her/his sixties’, and ‘woman’s/man’s face in her/his eighties’. Then, we use a pre-trained DEX to predict the age value of the image in

<https://github.com/siriusdemon/pytorch-DEX>

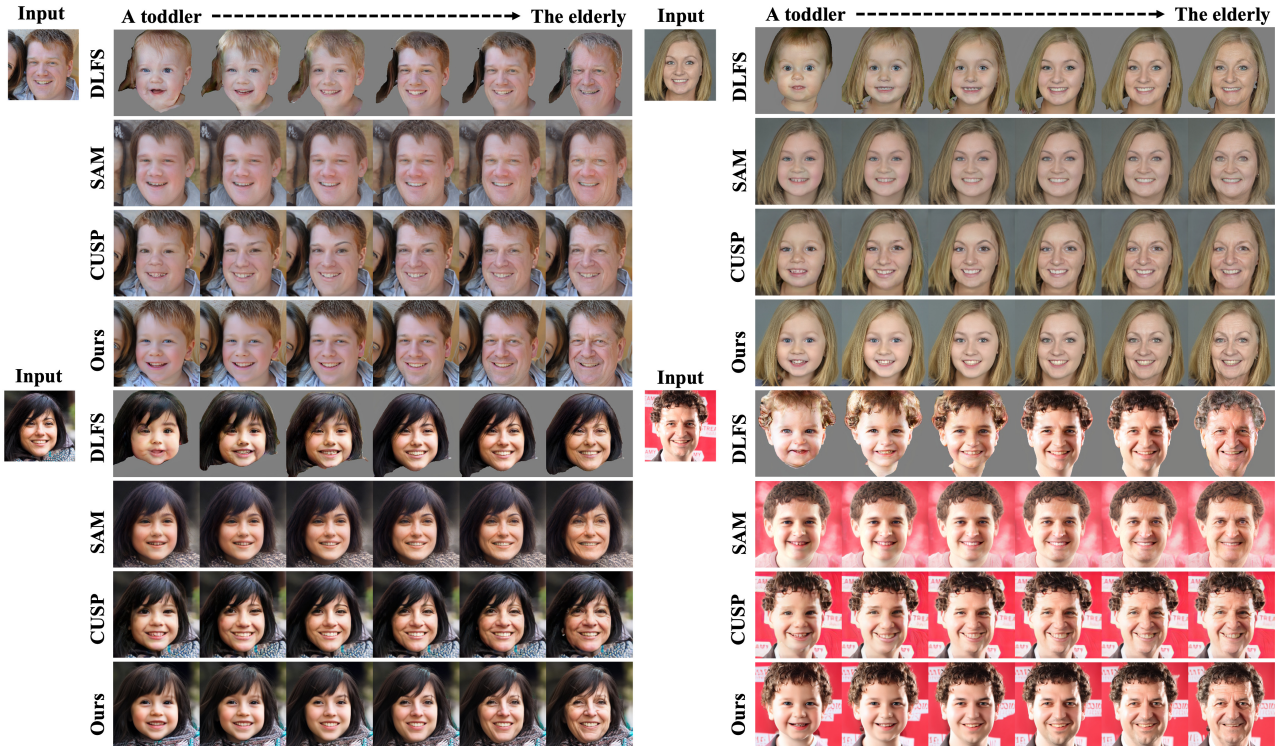


Figure 4: Qualitative comparison with DLFS [10], SAM [1], and CUSP [8] on the FFHQ-AT test set. Best viewed zoomed-in.

FFHQ. According to the absolute difference value between the predicted age and the pre-defined age, we annotate the image with the pre-defined text description. Note that, with these simple text descriptions, our PADA can achieve face aging based on arbitrary age-related descriptions, such as ‘*a face of teenager*’. We choose 66,928 images as training set and 3,072 images as test set, which contains 1,528 males and 1,544 females. Finally, we train our model on the new relabeled FFHQ-AT and evaluate it on both FFHQ-AT and CelebA-HQ [19] test sets.

Implementation Details. The conditional DDIM decoder and semantic encoder [31] are pre-trained on the FFHQ dataset [12]. In our L_{rec} , we randomly set the reference images equal to the source images with the probability of 0.167. Aging results are generated with $T = 25$ steps in all cases. Our implementations are based on the MindSpore. During training, we choose our hyperparameters with $\lambda_1 = 0.6$, $\lambda_2 = 0.01$, $\lambda_3 = 0.2$, $\lambda_4 = 0.01$, and $\lambda_5 = 0.1$. We set the margin m in L_{age} as 0.15 and the sampling intensity ϵ as 0.001, respectively. The Adam optimizer [15] is used with $lr = 0.0001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$. PADA is trained for 20 epochs with a batch size of 24.

Evaluation Metrics. We automatically and manually compare our method with three state-of-the-art aging methods, including DLFS[10], SAM[1], and CUSP[8]. We

compare them from three aspects: aging accuracy, identity preservation, and aging quality. For automatically evaluation, 1) since the pre-trained DEX [34] is used for training, we employ Face++ API for aging accuracy (Age MAE) evaluation; 2) since the pre-trained Arcface [5] is used for training, we employ the pre-trained SFace [39] for identity preservation evaluation; 3) for the aging quality evaluation, we employ Fréchet Inception Distance (FID) [11] to assess the discrepancy between the generated images and the real ones with the same age. For manual evaluation, we perform a human evaluation to reliably compare the performance with different aging methods [10, 1, 8]. Besides, we conduct the ablation study and explore more interesting properties of PADA, including diversity exploration, text-guided, and reference-guided face aging.

4.1. Comparison with Face Aging Methods

We compare our PADA with three state-of-the-art face aging methods: DLFS [10], SAM [1], and CUSP [8]. For all these methods, we employ their official implementation and pre-trained models. Note that the above methods are capable of covering different ranges of aging. Therefore, following [8], we compare the age group-based generation results with the three state-of-the-art methods.

Face++ Face detection API: <https://www.faceplusplus.com>.

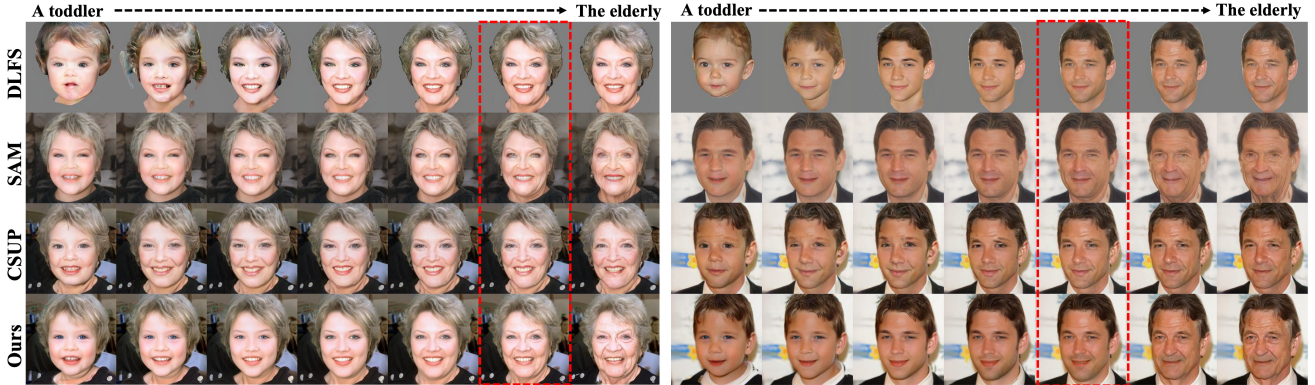


Figure 5: Generalization ability and comparison with DLFS [10], SAM [1], and CUSP [8] on the CelebA-HQ test set.

Qualitative Comparison. We show the face aging results from 0 to 70 years old with 10-20 age intervals for qualitative comparison on FFHQ-AT test set. Following DLFS [10], we generate missing aging images by interpolation in the latent space. As shown in Fig. 4, our PADA outperforms the three state-of-the-art face aging methods DLFS [10], SAM [1] and CUSP [8] in terms of shape deformation, texture transformation and generation quality. For shape deformation, our PADA and DLFS [10] can generate plausible baby faces with round shape and short baby teeth, while SAM [1] and CUSP [8] fail at shape deformation modeling, especially for babies. However, DLFS [10] tends to produce pronounced artifacts, resulting in less reliability. For texture transformation, our PADA and SAM [1] can produce detailed textures and achieve neck aging. CUSP [8] and DLFS [10] may generate artifacts on the teeth and facial contours. For generation quality, both our PADA and CUSP [8] can preserve pretty background information. Meanwhile, due to the GAN inversion [33], SAM [1] cannot accurately project real images into latent space, leading to blurred background. DLFS [10] can only generate facial images without background.

We also compare with the three state-of-the-art face aging methods DLFS [10], SAM [1], and CUSP [8] on CelebA-HQ [19] test set in Fig. 5. Obviously, our PADA outperforms these methods on both shape deformation and texture transformation. Besides, our PADA can generate geriatric spots (the right subject in Fig. 5), which cannot be achieved by other methods.

We also compare our PADA with DiffAE [31], which is a strong diffusion-based model for image manipulation. Here, we employ the official implementation of DiffAE [31] as baseline for face aging. As shown in Fig. 6, since DiffAE [31] (z^{stc} manipulation) achieves face aging by linear interpolation in the latent space, it ignores the non-linearity of aging process and cannot generate images of specified

Table 1: Quantitative comparison on FFHQ-AT test set.

	Age MAE (\downarrow)	ID Preservation (\uparrow)	FID (\downarrow)
DLFS[10]	10.36	0.5421	73.98
SAM[1]	9.31	0.5161	56.86
CUSP[8]	12.3	0.5320	37.01
Ours	9.19	0.6516	16.71

Table 2: Human evaluation on FFHQ-AT test set. Overall represents the results of all three age groups.

	DLFS [10]	SAM [1]	CUSP [8]	Ours	
Aging Accuracy (\uparrow)	Overall	32.33	38.87	37.86	46.29
	'a toddler'	46.67	37.05	32.95	46.41
	'an adult'	24.87	41.28	41.15	52.31
	'the elderly'	24.62	41.28	44.10	50.64
ID Preservation (\uparrow)	Overall	34.91	40.68	32.78	40.80
	'a toddler'	56.05	57.41	54.15	56.73
	'an adult'	54.97	55.10	54.28	56.19
	'the elderly'	57.73	59.87	58.40	60.53
Aging Quality (\uparrow)	Overall	36.80	44.15	40.18	55.54

ages. PADA (z^{age} manipulation) achieves more reliable aging results with complex non-linear aging directions. Moreover, we incorporate CLIP encoders for flexible face aging and PAE for pluralistic aging with high-level variances.

Quantitative Comparison. Both aging accuracy and identity preservation are essential quantitative metrics for face aging. We report the quantitative comparison results with DLFS [10], SAM [1], and CUSP [8] in Table 1. As expected, our PADA achieves the best performance on both the identity preservation and aging accuracy (Age MAE). Besides, our PADA also achieves best performance on aging quality (FID [11]), which supports the qualitative analysis in Fig. 4 and Fig. 5.

Human Evaluation. To provide more reliable quantitative analysis, we perform a human evaluation to compare different methods. Following SAM [1], we compare different methods on three age groups: 'a toddler', 'an adult',

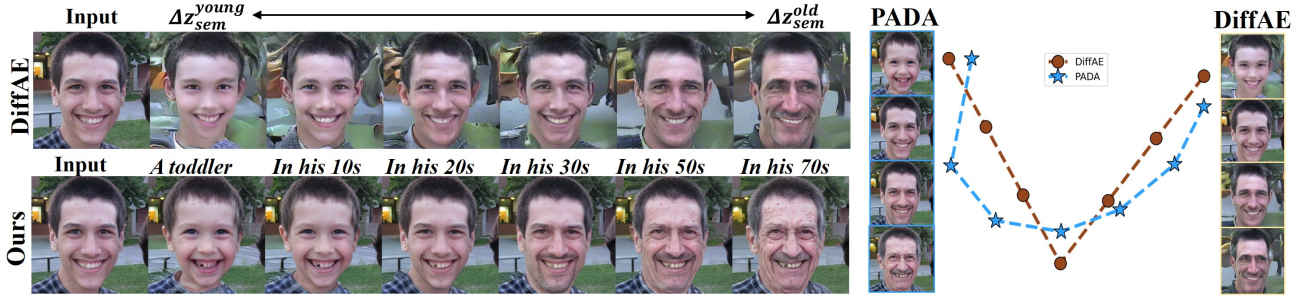


Figure 6: Compared with DiffAE. To compare the learned aging directions of DiffAE and PADA, we use PCA to project the learned latent codes obtained from 14 different images (right). Zoomed-in.




Figure 7: Pluralistic face aging results of PADA. The first row shows the diverse aging results with high-level variations. The second row shows the results with low-level variations.

and ‘the elderly’. Specifically, we invite 60 volunteers and ask them to score the results of different methods, according to aging accuracy, identity preservation and aging quality. Each volunteer is randomly allocated 30 reference image sets. Note that all images are randomly selected. We provide the human evaluation results in Table 2. Overall represents the results of all three age groups. As expected, our PADA consistently performs best in all three metrics. We observe that the aging accuracy of DLFS[10] on ‘a toddler’ is a little better than ours, while the identity preservation of SAM[1] on ‘a toddler’ is a little better than ours. This is because the aging accuracy and identity preservation may conflict when we translate an adult into a toddler.

4.2. Ablation Study

Diversity Exploration. Due to the sequential denoising reverse process and the probabilistic aging embedding learning, our PADA generates pluralistic aging results with both low-level variations and high-level variations. We show the diverse aging results in Fig. 7. The high-level variations are relevant to face shape and skin color, while the low-level variations are relevant to the types and locations of wrinkles.

Text-guided and Reference-guided Face Aging. Thanks to our strategy of probabilistic aging embedding

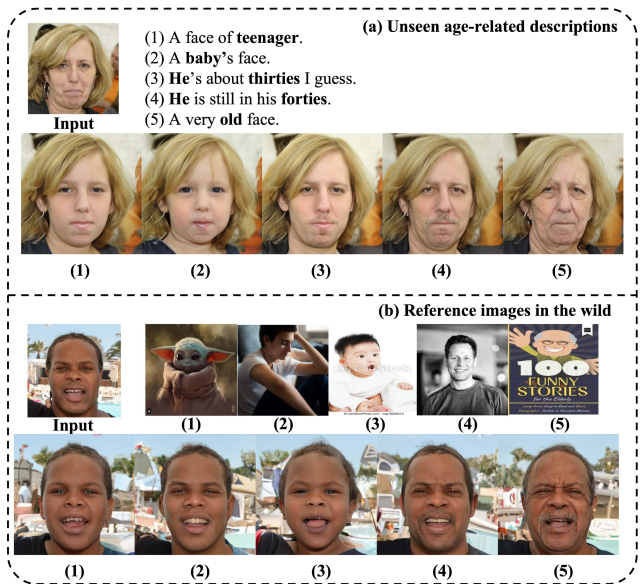


Figure 8: Face aging conditioned on unseen age-related descriptions and reference images in the wild. (a) Despite never being trained with texts of ‘a very old face’, our PADA still yields plausible face aging results. (b) We can utilize arbitrary reference images to guide the aging process.

learning in CLIP latent space, our PADA can achieve face aging conditioned on the unseen age-related text descriptions or arbitrary face images from the Internet. Thus, both the text-driven and reference-driven interfaces for face aging are provided. The results based on the open-world age descriptions or arbitrary unseen facial images are shown in Fig. 8. Although our PADA has not seen both these two variants during training, it still can generate plausible face aging results.

The Effectiveness of Losses. We report the qualitative visualization results in Fig. 9 for a comprehensive comparison between our PADA and its five variants. The quantitative comparison is included in supplementary materials. As

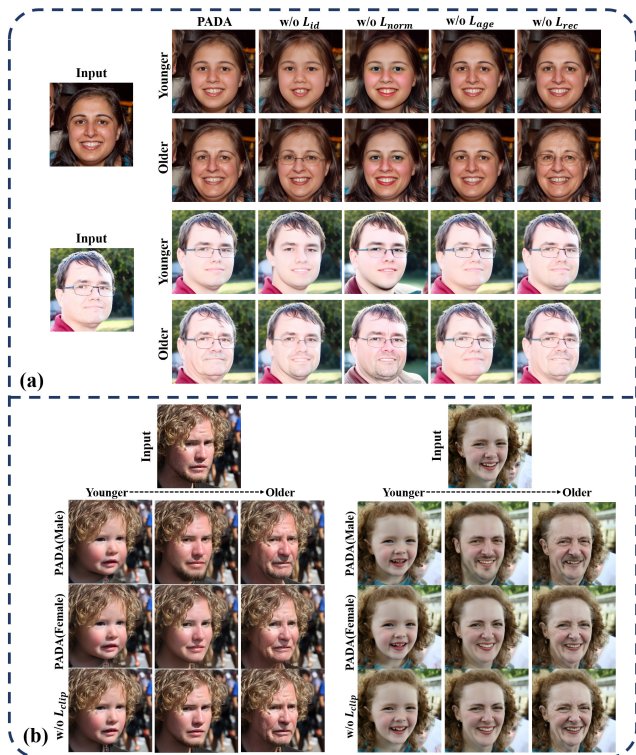


Figure 9: Visual comparisons of our method with its variants.

shown in Fig. 9 (a), without L_{id} loss, much identity information is lost during aging process. The pose and expression information are also dramatically destroyed. Without L_{norm} loss, the generation quality is degraded and unrealistic images with artifacts are produced, e.g., the blush on the face. The lack of L_{age} leads to almost negligible aging changes. L_{rec} is randomly applied for each batch during training, but it can preserve age-irrelevant information.

As shown in Fig. 9 (b), L_{clip} improves the aging fidelity from two perspectives. (I) It makes our PADA learn more aging details from the pre-trained CLIP, e.g., the beards and necklines. Additionally, it leads to better shape transformation for babies; (II) It can accurately model aging process for different genders. Without L_{clip} , our PADA can hardly model the clearly gender-specific aging process. These indicate that each component of PADA is essential.

Table 3 further presents the performance on aging accuracy and identity preservation of different variants of our PADA on FFHQ-AT test set. As expected, without the preservation loss: L_{id} , L_{norm} , or L_{rec} , the performance of identity preservation decreases. Without the age fidelity loss: L_{age} or L_{clip} , the performance of aging accuracy has dropped. These indicate that each component in our method is essential for synthesizing photo-realistic aging results.

Table 3: Model comparisons on FFHQ-AT.

Method	Age MAE (\downarrow)	ID Preservation (\uparrow)
w/o L_{id}	8.70	0.4766
w/o L_{norm}	9.23	0.6446
w/o L_{rec}	9.28	0.4516
w/o L_{clip}	11.07	0.7046
w/o L_{age}	16.62	0.7176
w/o L_{kl}	9.99	0.6736
Ours	9.19	0.6516

Additionally, we observe that without the age fidelity loss: L_{age} or L_{clip} , the model can achieve better identity preservation. Meanwhile, without the preservation loss: L_{id} or L_{norm} , the performance of aging accuracy is improved. All of the above phenomena are reasonable. This is because that there may be conflict between aging accuracy and identity preservation. Specifically, according to the survey of age-invariant face recognition [29], both the shape and texture changes degrade the performance of face recognition systems. The employ of age fidelity loss leads to the shape and texture changes, which degrades the performance of identity preservation. While the employ of the preservation loss inhibits the shape or texture changes, which degrades the performance of aging accuracy. Our PADA better balances the aging accuracy and identity preservation, achieving plausible face aging with pretty identity preservation.

5. Conclusions

In this paper, we have proposed a CLIP-driven Pluralistic Aging Diffusion Autoencoder (PADA) to achieve diverse plausible face aging with both low-level stochastic variations and high-level aging semantic variations. To produce stochastic low-level aging details, we resort to diffusion models by sequential denoising reverse. To learn diverse high-level aging patterns, we present Probabilistic Aging Embedding (PAE) in the common CLIP latent space, where a text-guided KL-divergence loss is imposed to guide the learning of distribution. Our PADA can achieve pluralistic face aging conditioned on both the open-world age-related descriptions and arbitrary unseen facial images. Extensive experiments demonstrate that our method obtains state-of-the-art face aging results in terms of generation quality and diversity.

Acknowledgment This work is sponsored by Beijing Nova Program (Z211100002121106), National Natural Science Foundation of China (Grant No.62006228, Grant No. 62176025), Youth Innovation Promotion Association CAS (Grant No.2022132), and CAAI-Huawei MindSpore Open Fund.

References

- [1] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Only a matter of style: Age transformation using a style-based regression model. *TOG*, 40(4):1–12, 2021.
- [2] Fei-Long Chen, Du-Zhen Zhang, Ming-Lun Han, Xiu-Yi Chen, Jing Shi, Shuang Xu, and Bo Xu. Vlp: A survey on vision-language pre-training. *Machine Intelligence Research*, 2023.
- [3] Sanghyuk Chun, Seong Joon Oh, Rafael Sampaio De Rezende, Yannis Kalantidis, and Diane Larlus. Probabilistic embeddings for cross-modal retrieval. In *CVPR*, pages 8415–8424, 2021.
- [4] Biplob Deb Nath, Giuseppe Coviello, Yi Yang, and Srimat Chakradhar. Uac: An uncertainty-aware face clustering algorithm. In *ICCV*, pages 3487–3495, 2021.
- [5] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, pages 4690–4699, 2019.
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 34, 2021.
- [7] Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *TOG*, 41(4):1–13, 2022.
- [8] Guillermo Gomez-Trenado, Stéphane Lathuilière, Pablo Mesejo, and Óscar Cerdón. Custom structure preservation in face aging. In *ECCV*. Springer, 2022.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014.
- [10] Sen He, Wentong Liao, Michael Ying Yang, Yi-Zhe Song, Bodo Rosenhahn, and Tao Xiang. Disentangled lifespan face synthesis. In *ICCV*, pages 3877–3886, 2021.
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30, 2017.
- [12] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019.
- [13] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, pages 8110–8119, 2020.
- [14] Gwanghyun Kim and Jong Chul Ye. Diffusionclip: Text-guided image manipulation using diffusion models. In *CVPR*, 2022.
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *arXiv preprint arXiv:2107.00630*, 2021.
- [17] Diederik P Kingma and Max Welling. An introduction to variational autoencoders. *arXiv preprint arXiv:1906.02691*, 2019.
- [18] Gihyun Kwon and Jong Chul Ye. Clipstyler: Image style transfer with a single text condition. In *CVPR*, pages 18062–18071, 2022.
- [19] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *CVPR*, pages 5549–5558, 2020.
- [20] Gun-Hee Lee and Seong-Whan Lee. Uncertainty-aware mesh decoder for high fidelity 3d face reconstruction. In *CVPR*, pages 6100–6109, 2020.
- [21] Peipei Li, Yibo Hu, Ran He, and Zhenan Sun. Global and local consistent wavelet-domain age synthesis. *IEEE TIFS*, 14(11):2943–2957, 2019.
- [22] Peipei Li, Yibo Hu, Qi Li, Ran He, and Zhenan Sun. Global and local consistent age generative adversarial networks. In *ICPR*, pages 1073–1078. IEEE, 2018.
- [23] Peipei Li, Yibo Hu, Xiang Wu, Ran He, and Zhenan Sun. Deep label refinement for age estimation. *PR*, 100:107178, 2020.
- [24] Peipei Li, Huaibo Huang, Yibo Hu, Xiang Wu, Ran He, and Zhenan Sun. Hierarchical face aging through disentangled latent characteristics. In *ECCV*. Springer, 2020.
- [25] Yunfan Liu, Qi Li, and Zhenan Sun. Attribute-aware face aging with wavelet-based generative adversarial networks. In *CVPR*, pages 11877–11886, 2019.
- [26] Farkhod Makhmudkhujaev, Sungeun Hong, and In Kyu Park. Re-aging gan: Toward personalized face age transformation. In *ICCV*, pages 3908–3917, 2021.
- [27] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, pages 8162–8171. PMLR, 2021.
- [28] Roy Or-El, Soumyadip Sengupta, Ohad Fried, Eli Shechtman, and Ira Kemelmacher-Shlizerman. Lifespan age transformation synthesis. In *ECCV*, pages 739–755. Springer, 2020.
- [29] Unsang Park, Yiying Tong, and Anil K Jain. Age-invariant face recognition. *IEEE TPAMI*, 32(5):947–954, 2010.
- [30] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *ICCV*, pages 2085–2094, 2021.
- [31] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizatwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *CVPR*, pages 10619–10629, 2022.
- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021.
- [33] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *CVPR*, pages 2287–2296, 2021.
- [34] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *IJCV*, 126(2):144–157, 2018.
- [35] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2020.

- [36] Jianxin Sun, Qiyao Deng, Qi Li, Muye Sun, Min Ren, and Zhenan Sun. Anyface: Free-style text-to-face synthesis and manipulation. In *CVPR*, pages 18687–18696, 2022.
- [37] Tianyi Wei, Dongdong Chen, Wenbo Zhou, Jing Liao, Zhen-tao Tan, Lu Yuan, Weiming Zhang, and Nenghai Yu. Hair-clip: Design your hair by text and reference image. *arXiv preprint arXiv:2112.05142*, 2021.
- [38] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *CVPR*, pages 5810–5818, 2017.
- [39] Yaoyao Zhong, Weihong Deng, Jiani Hu, Dongyue Zhao, Xian Li, and Dongchao Wen. Sface: Sigmoid-constrained hypersphere loss for robust face recognition. *IEEE TIP*, 30:2587–2598, 2021.