

Identity-Preserving Aging and De-Aging of Faces in the StyleGAN Latent Space

Luis S. Luevano Pavel Korshunov Sébastien Marcel

Idiap Research Institute

Rue Marconi 19, 1920 Martigny, CH

{luis.luevano, pavel.korshunov, sebastien.marcel}@idiap.ch

Abstract

Face aging or de-aging with generative AI has gained significant attention for its applications in such fields like forensics, security, and media. However, most state of the art methods rely on conditional Generative Adversarial Networks (GANs), Diffusion-based models, or Visual Language Models (VLMs) to age or de-age faces based on pre-defined age categories and conditioning via loss functions, fine-tuning, or text prompts. The reliance on such conditioning leads to complex training requirements, increased data needs, and challenges in generating consistent results. Additionally, identity preservation is rarely taken into account or evaluated on a single face recognition system without any control or guarantees on whether identity would be preserved in a generated aged/de-aged face. In this paper, we propose to synthesize aged and de-aged faces via editing latent space of StyleGAN2 using a simple support vector modeling of aging/de-aging direction and several feature selection approaches. By using two state-of-the-art face recognition systems, we empirically find the identity preserving subspace within the StyleGAN2 latent space, so that an apparent age of a given face can be changed while preserving the identity. We then propose a simple yet practical formula for estimating the limits on aging/de-aging parameters that ensures identity preservation for a given input face. Using our method and estimated parameters we have generated a public dataset of synthetic faces at different ages that can be used for benchmarking cross-age face recognition, age assurance systems, or systems for detection of synthetic images. Our code and dataset are available at the project page <https://www.idiap.ch/paper/agesynth/>

1. Introduction

Face image synthesis has seen substantial advancements in recent years, driven by breakthroughs in generative models like Generative Adversarial Networks (GANs) [16, 21]. Applications for face image synthesis scenarios include face aging and de-aging, strengthening face [28] and age [26]




Figure 1. Samples from fully synthetic dataset generated. IDs generated using Langevin sampling [13] and linear approach for aging and de-aging the synthetic identities.

verification systems against physical [24, 30] and digital [23, 29] spoofing attacks, mitigating bias of face recognition (FR) systems across different age groups [25], and improving robustness of cross-age face recognition. In particular, these methods can be used to modify a person’s face appearance to an arbitrary age while retaining their identity, allowing us to generate younger and older looking versions of the same subject. State of the art approaches typically employ Conditional GANs [3] or Diffusion models [5] that condition the face transformation on a specific target age group, producing realistic results at the expense of sacrificing controllability, involving complex training procedures, extensive computing resources, and requiring large amounts of labeled data for fine tuning.

When faces are synthetically projected aged or de-aged, it is difficult to preserve identity, which is essential in applications where age transformation must retain enough distinctive features for accurate recognition purposes. While some face aging methods incorporate identity consistency

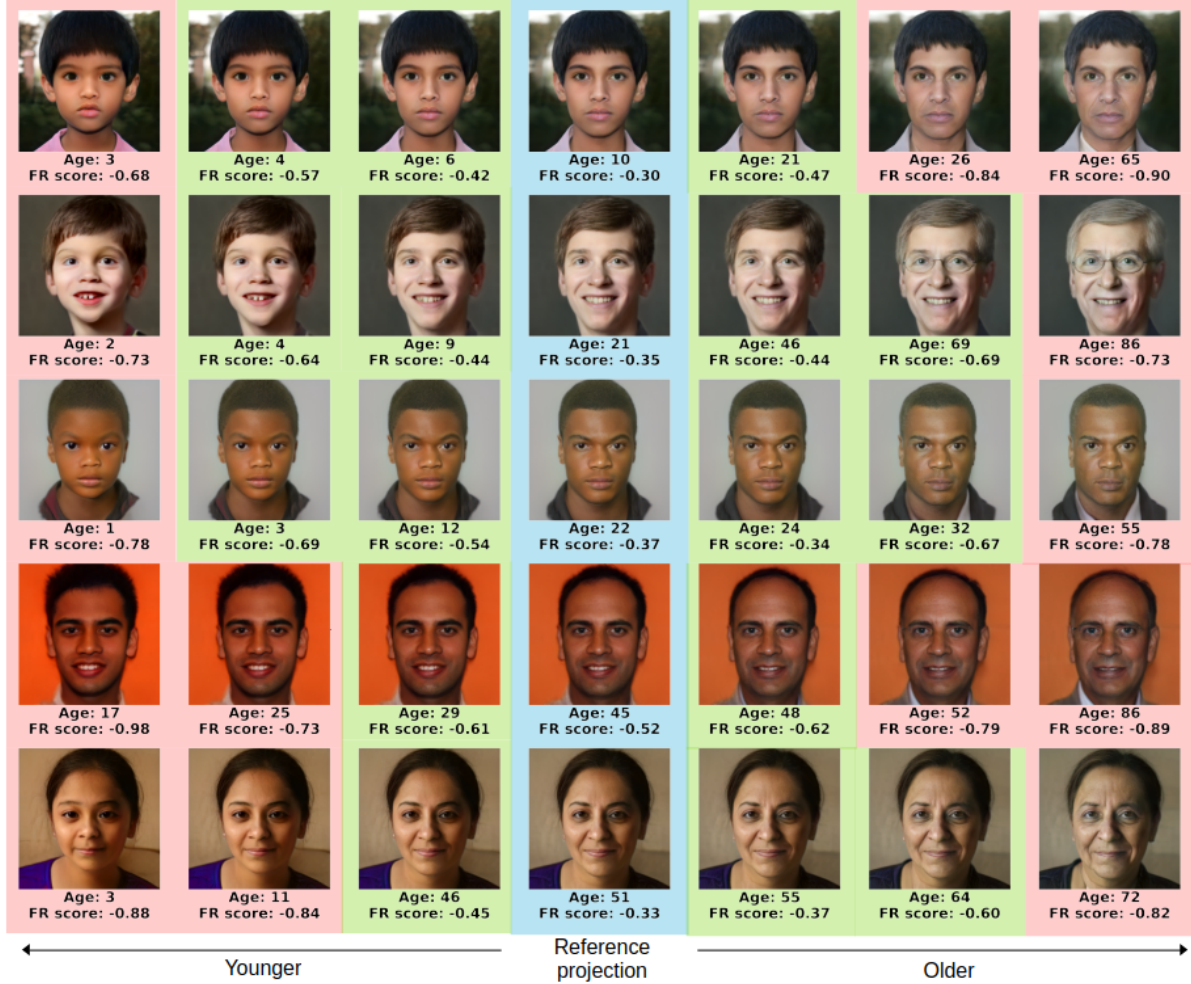


Figure 2. UTKFace sample re-projections from different age groups showing variation in estimated age and FR score from the cosine distance of the EdgeFace embedding to the real-world reference. The column with blue background shows the re-projection without aging/de-aging, the aged/de-aged faces over green are verified against the original, while those over red are rejected.

losses [34] or use pre-trained face recognition models [6] to verify the synthesized face’s identity, most approaches are limited by their dependence on a single face recognition model feature space when conditioning GANs [3, 4] and posterior face verification [34]. This reliance on using existing face recognition pretrained models can potentially introduce biases, as the preservation of identity may vary when evaluated using different recognition systems. As such, it is important to include evaluations using different face recognition systems, whereas current age synthesis methods rarely report face recognition accuracy performance on more than one system.

In particular, GAN approaches operate in continuous latent spaces and lack predefined steps or directions to represent a person’s age, making them challenging to navigate effectively for cross-age synthesis while preserving identity. In this work, we address this gap by conducting targeted experiments in the StyleGAN2 [22] latent space, a GAN vari-

ant popular for its ability to represent style-based attributes, to identify a global “age synthesis direction”, examining the extent to which identity can be preserved across different age groups. We further enhance identity retention by selecting relevant features, both with and without explicit age or face recognition labels. Furthermore, we generate and release a fully synthetic dataset using face age synthesis approach and synthetic identity sampling, illustrated in Figure 1.

Concretely, in this work we review the current state of face age synthesis research and propose a simple data-efficient strategy that operates within the StyleGAN2 W latent space. We explore the challenges of navigating this space while maintaining robust face verification performance (as shown in Figure 2) using a linear approach for age synthesis and improve identity retention by employing various feature selection mechanisms, including comparisons using reconstructions from Principal Component

Analysis (PCA) and Linear Discriminant Analysis (LDA) subspaces given the original projected datasets in the latent space. Unlike conventional methods that rely on conditioning GANs, our approach leverages the readily-available StyleGAN2 structured and rich latent space for straightforward edits that achieve both aging and de-aging without re-training GAN encoders and decoders. We assess how different feature selection approaches, particularly those emphasizing age and identity components, affect identity preservation when using the selected components in latent space manipulations. We then map the relationship between apparent age, StyleGAN2 scalar step, and assess the identity preservation bounds via curve fitting and root analysis to propose a more practical approach for synthesizing specific ages using this latent space. Finally, we generate a fully synthetic public dataset¹ with aged face images by applying the aging approach on synthetically generated identities with Langevin sampling [13].

This paper has the following contributions:

1. We provide a baseline streamlined and data-efficient approach to face age synthesis using latent editing within StyleGAN2, reducing the need for conditioning networks and extensive labeled datasets by extending linear approaches using Support Vector Regression.
2. We propose different feature selection strategies and evaluate them using the latents from the StyleGAN space, including reconstruction approaches based on Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA), that effectively improve identity preservation and apparent age synthesis within the StyleGAN2 latent space.
3. We propose a set of specific experiments and visualization techniques to analyze the extent of identity preservation within the StyleGAN2 latent space across age modifications, identifying thresholds where identity starts to degrade.
4. We release a practical toolset¹ for synthesizing faces to a specific age with our findings: linear age direction, weights of W components, fitted curves per age group, and code for finding the corresponding StyleGAN2 scalar step.
5. We generate and release a synthetically generated dataset¹ using our approach with 20k identities and 10 age variations per identity, providing replicability, and addressing privacy concerns over the sensitivity of cross-age data availability.

The paper is organized as follows: Section 2 describes the related work for face age synthesis methods, Section 3

details approach for efficient controlled latent editing and feature selection, Section 4 shows our experimental setup, in Section 5 we present the experiment results, and we conclude in Section 6.

2. Related work

In this section we present related methods on face age synthesis, for Deep Learning-based approaches. It is possible to classify these methods on those based on Translation, Sequences, Conditions, and Style [1, 20].

Translation-based methods. These approaches use GAN methods to age faces by training generators to address consistency among specific age gaps. The authors of F-GAN [33] propose to use CycleGAN [39] losses and face age estimators to improve synthesis between smaller gaps in age groups. E2E-CycleGAN [36] trains using edge maps and generates them before synthesizing the textures, helping preserve geometry. FA-CGAN [4] uses the CycleGAN backbone and with additional convolutional blocks, an alternative discriminator, and multi-scale feature fusion to improve face image quality including fine details such as wrinkles.

These approaches show limitations because they must train multiple generator–discriminators with pairs for every age gap, while relying on cycle and identity losses that depend on face recognition backbones. They also rely on training specific pairs of generators for aging and de-aging predetermined age gaps, becoming computationally expensive to age or de-age outside of these age groups, requiring additional generators.

Sequence-based methods. These methods decompose the age gaps into a sequence of incremental aging steps. Recurrent Face Aging (RFA) [35] uses optical–flow-normalized eigenfaces as input and trains Gated Recurrent Unit (GRUs) to model long and short-term dynamics of aging, transferring high-frequency details from each group’s nearest neighbors. Triple-GAN [12] aligns three simultaneous translation paths with a “triple” loss so that the generated faces along the different paths converge to the same target age. PFA-GAN [19] divides the generator itself into gated sub-networks, each tuned to a narrow age span, yielding smoother trajectories and mitigating ghosting artifacts.

Shortcomings of these approaches include cases when we attempt to synthesize over large age gaps, such as error accumulation and an increase in complexity due to the chain of generators. Furthermore, it is not possible to synthesize other age gaps different from those contemplated at training time.

¹Code and dataset available at the project page:
<https://www.idiap.ch/paper/agesynth/>

Conditional-based methods. Conditional GANs use the desired age label as input for the generator and discriminator. Age-CGAN [3] searches the latent space with identity-preserving optimisation so that FaceNet embeddings stay close while the age code changes. IPC-GAN [34] adds perceptual, age-classification, and identity preservation losses at training time. Other methods, such as ChildFace [8], AW-GAN [6], and AgeTransGAN [18] strengthen detail using coarse-to-fine generators, wavelet-domain critics, and multitask discriminators, respectively. In ChildGAN [7], the authors train separate StyleGAN2 instances per gender and perform logistic regression in the latent space for attribute augmentation.

These methods show limitations in their dependence on explicit target age labels at train time and test time. They also need large and accurately annotated age datasets and employ complex conditioning pipelines for synthesizing multiple attributes such as age, gender, or pose. Their identity preservation mechanisms also rely heavily on a single recognition model embedded in the loss, potentially limiting the usability of the output to a particular face recognition backbone.

Style-based methods. These approaches shift the focus from re-training to navigating an already expressive latent space. SAM [2] places a head for age regression on top of StyleGAN2 and back propagates the regression gradient to the W projection modules, regularized by identity and cycle constraints. FAM [10] directly manipulates FaceNet features before re-synthesizing the image. Other methods include IricGAN [32], which uses multi-scale attribute regression, and DyStyle [27], comprised of dynamic expert modules selected per attribute. HDA-SynChildFaces [11] employs multiple SVM hyper-planes in the latent space to generate a large dataset focused on children faces only. Diffusion models for aging faces, such as [5], use text prompts and labels as input to the method, delegating the text-to-image translation task to cross-modal alignment methods, thus sacrificing controllability.

Even though these approaches are based on latent editing, the training process is often made more complex by re-training the generator, adding extra networks at training time, introducing complex training losses, and finding latent directions using *multiple* SVMs. Furthermore, these approaches do not attempt to explore which latent components meaningfully affect age or identity, which may also introduce artifacts and undesired identity changes. In this paper, we address these shortcomings by using the readily-available StyleGAN2 generator without re-training or training additional convolutional modules, finding a *single* latent direction with little amount of training data, and improving identity preservation by selecting relevant latent vector components to age and identity.

3. Editing in the StyleGAN latent space

In this section, we will detail the method for latent analysis, latent edits, and feature selection.

3.1. Latent analysis and linear latent editing

The latent analysis approach is an extension of the method from [9], which begins by projecting real face images from the image space into the W 512- d dimensional latent space of StyleGAN2, where the w latent vectors are used for age attribute analysis and editing. We extend the approach in [9] for finding the direction in the latent space from ‘young’ to ‘old’. Our approach uses a Linear SVR to fit latents in the StyleGAN2 W space to their original age labels $y \in \mathbb{R}$. Linear SVR fits the hyperplane containing most of the samples inside its surrounding margin. This hyperplane is defined by $\{y = \theta^T \tilde{w} | \theta, \tilde{w} \in \mathbb{R}^{|W|+1}\}$, where θ contains the hyperplane parameters (bias b and coefficients λ) and \tilde{w} a latent vector $[1 \ w]$ in $\mathbb{R}^{|W|+1}$. Given that we have a known latent w_0 with age value y_0 , we want to find a new latent w_1 reflecting the desired change in age (we want to age or de-age w_0). By projecting the vector θ , used to predict the y age value, onto the latent vector \tilde{w}_0 , we can find the direction in which the age of w_0 changes within W space. We can formulate $\tilde{w}_1 = \tilde{w}_0 + c\tilde{w}_0$, with c given by $\text{Proj}_{\tilde{w}_0} \theta = \frac{\theta \cdot \tilde{w}_0}{\tilde{w}_0 \cdot \tilde{w}_0}$. After substitution, we have \tilde{w}_1 as :

$$\tilde{w}_1 = \tilde{w}_0 + \theta$$

$$[1 \ w_1] = [1 \ w_0] + [b \ \lambda]$$

With λ corresponding to the vector changing the components of w_0 in the W space. Finding the direction of λ as $\hat{\lambda} = \frac{\lambda}{\|\lambda\|_2}$, a normalized unit vector, allows us to move w_0 to the desired w_1 by scaling $\hat{\lambda}$ using scalar multiplication s as:

$$w_1 = w_0 + s\hat{\lambda} \quad (1)$$

By projecting w_1 into pixel space, we obtain a face of the person appearing older or younger depending on the sign of s , which its value represents the number of years. What remains to be found is the mapping between the values of scalar s and the number of years by which we want to age/de-age the person. We are interested in mapping the changes in age in function of s because obtaining w_1 from w_0 , the y value in years, and the age direction $\hat{\lambda}$ depends on moving through a linear hyperplane, whereas the latents may not correspond linearly to their y age value and their placement in the latent space may be depending on other factors, such as the apparent age of the original face.

3.2. Identity preservation in the latent space

When using Eq. 1, we move linearly using all the components of the StyleGAN2 vector in the latent space indiscriminately by multiplying them with the floating-point scalar

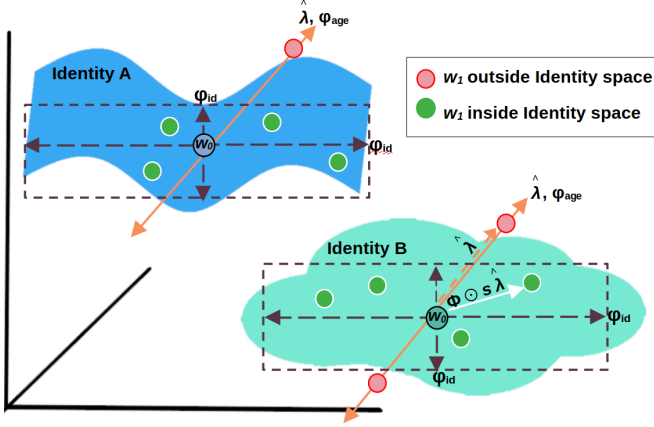


Figure 3. Controlled latent edit approach with identity preservation in the W space. We age and de-age each sample projection w using the scaled age direction $s\hat{\lambda}$, controlled by the rate weight Φ , which includes the change in the selected latent vector components relevant to identity (φ_{id}) and age (φ_{age}). This re-scales the change in the age direction, aiding in synthesizing points inside the identity region of each subject inside the latent space.

step s . We refer to the approach described by Eq. 1 as a linear baseline approach, since it does not involve any identity constraints.

We can extend the linear baseline approach by multiplying only a selected subset of the components in the StyleGAN2 vector, with the aim to find those components that contribute the most to the identity and age changes. Hence, we can modify Eq. 1 to include a Φ term describing an additional weight vector regulating the change per each component of the scaled $\hat{\lambda}$ age direction vector as follows:

$$w_1 = w_0 + \Phi \odot s\hat{\lambda} \quad (2)$$

The Φ vector re-scales and re-shapes the region followed by the synthesized latent vector w_1 through element-wise multiplication (\odot), constraining the magnitude of the directions for different components of the latent vector. We compute this Φ vector from selecting features relevant to age (φ_{age}) and identity (φ_{id}) from the latent components. In the linear baseline case, this Φ scaling vector represents an equal change in all the components from $\hat{\lambda}$ and w . Figure 3 illustrates this approach.

Feature selection. The main goal of feature selection in this work is to compute Φ from Eq. 2 to regulate the change in our latent vector components. We propose multiple ways of computing Φ . In general, we employ PCA and LDA to assess component importance through sub-space projection and reconstruction. For this purpose, we use two sets of *standardized* latent vectors: one from a dataset with identity labels ($V_{id}^{n \times |W|}$) (for PCA and LDA) and another one from a dataset with age group labels ($V_{age}^{m \times |W|}$), (only for LDA).

Computing Φ using PCA. We use the V_{id} set to compute $\Phi_{PCA} = \varphi_{id-PCA}$. Firstly, we compute the eigen-decomposition for $aP_{PCA} = \text{Eig}(V_{id}V_{id}^T)$, with a denoting the vector of unsorted eigenvalues of the projection basis P_{PCA} . We calculate the cumulative variance of the largest eigenvalues and select the component indices corresponding to 95% of variance (σ_{PCA}). We use the selected components to directly set the values of the φ_{id-PCA} mask, since we found that the most relevant components in the PCA basis are the same as in the latent space after projection, reconstruction, and comparison. Consequently, we formulate:

$$\varphi_{id-PCA} = \{\rho_i \dots \rho_{|W|} | \rho = 1 \text{ if } a_i \in \sigma_{PCA} > 95\%, 0 \text{ otherwise}\}$$

Computing Φ using LDA. We find the components important for identity φ_{id-LDA} using V_{id} with identity labels and those important for age as $\varphi_{age-LDA}$ with V_{age} with age group labels.

We compute separate LDA bases, aP_{LDA} , from the eigen-decomposition of the product of the scatter matrices for each set of latents. The magnitude of the eigenvalues in a represent the ratio of discriminability per class and component, which we use to find the most important discriminative components in the basis, leaving those contributing to 95% discriminability intact. We set the least important components in P_{LDA} as zero vectors, and use this matrix to project the latents into the sub-space. Using the matrix pseudoinverse P_{LDA}^+ , we reconstruct the projected latents from the LDA space back to the latent space to compare them against the original latents. We denote this reconstruction V^* .

Each pair of standardized latents V and reconstructed latents V^* is used to calculate three different distance metrics $\Psi(V, V^*)$: Mean Squared Error (MSE), 1-D Wasserstein distance, and Covariance. We compute the mean per component of $\Psi(V, V^*)$ as $\psi^{|W|}$. We calculate the mean μ_ψ as threshold for feature selection. We formulate φ_{LDA} when using MSE and 1-D Wasserstein as:

$$\varphi_{LDA} = \{\rho_i \dots \rho_{|W|} | \rho = 1 \text{ if } \psi_i < \mu_\psi, 0 \text{ otherwise}\}$$

and for Covariance as:

$$\varphi_{LDA} = \{\rho_i \dots \rho_{|W|} | \rho = 1 \text{ if } \psi_i > \mu_\psi, 0 \text{ otherwise}\}$$

We use these formulations to calculate the masks φ_{id^*-LDA} and φ_{age^*-LDA} corresponding to each different distance metric. These masks correspond to *all* the components important to identity or age. Hence, we find an overlap of the components relevant to both identity and age as:

$$\varphi_{id \wedge age} = \varphi_{id^*-LDA} \wedge \varphi_{age^*-LDA}$$

while the components relevant *only* to age and identity are formulated as

$$\varphi_{age-LDA} = \varphi_{age^*-LDA} \wedge \neg \varphi_{id^*-LDA}$$

$$\varphi_{\text{id-LDA}} = \varphi_{\text{id*} \text{-LDA}} \wedge \neg \varphi_{\text{*age-LDA}}$$

Therefore, we can transform the weighting vector Φ from Eq. (2) by including the features selected using LDA as follows:

$$\Phi_{\text{LDA}} = \alpha \varphi_{\text{age-LDA}} + \beta \varphi_{\text{id} \wedge \text{age}} \quad (3)$$

where the hyperparameters α and β are scalars emphasizing the specific components that are relevant only for age and those impacting age and identity at the same time.

When using LDA with the age-labeled dataset, we could have few age group labels, which could potentially limit the effect of LDA since it is more suitable for data with larger number of classes. In our experiments, we therefore applied LDA with larger number of age classes but did not observe any impact in the component selection results.

3.3. Estimating scalar steps for specific ages

When we follow Eq. (2) for aging and de-aging, we compute StyleGAN latents and the corresponding images corresponding to a discrete small number of scalars s for each age group. For practical application, we need to know scalar values for every desired age y_r we want to generate.

We then fit polynomial models $p_r(s)$, to approximate $p_r(s) \approx y_r$ based on the latents we pre-computed for the limited number of scalar s . When transforming a person from an *original* age y_{original} to a *desired* age y_{desired} , we select the polynomial curve p_r corresponding to the sample's age range r from y_{original} . To find the original and desired ages in the polynomial curve, s_{original} and s_{desired} respectively, we solve for the scalar shifts by computing the roots of $p_r(s) - y_{\text{original}} = 0$ and $p_r(s) - y_{\text{desired}} = 0$. The solutions correspond to the scalar value at that age, s_{original} and s_{desired} . We compute the net scalar offset as $\Delta s = s_{\text{desired}} - s_{\text{original}}$. In practice, we only keep physically meaningful solutions by filtering out those outside a functional scalar range. If a real and unique valid solution does not exist for computing the scalar using the polynomial curve in a category, we can fall back to simpler linear fits, separately for the aging and de-aging cases.

4. Experimental setup

In this section, we describe our experiments in terms of datasets, practical latent editing, features selection, and evaluation pipeline.

4.1. Datasets

In the experiments, we use UTKFace [37] for the age analysis and Color FERET [15] for the identity analysis.

UTKFace. This dataset [37] contains samples with age data ranging from 0 to 116 years old, with over 20K faces

annotated with age, gender, and ethnicity attributes. We use only a small portion of images to simulate data-constrained scenarios. Firstly, we pre-process the images from the UTKFace dataset using CodeFormer [38] to scale them up to $2\times$ their original resolution, cropped and aligned to match the original StyleGAN2 FFHQ images to a resolution of $1,024 \times 1,024$ pixels. Next, we take a section of the upscaled dataset, dividing it into a training set partition and a testing set partition. The training set consists of 1,336 samples, where half are younger than 18, and the other half are over 25. The test set contains 3,494 samples. We ensure no overlap with the training sets using an age estimation model [26]. We also group the samples in each set using a different number of age groups, which is useful depending on the experiment. For 9 age groups, we divide them for < 8 , $[8 - 13)$, $[13, 18)$, $[18 - 25)$, $[25 - 35)$, $[35 - 45)$, $[45 - 55)$, $[55 - 65)$, and > 65 . When using 4 age groups, we label the ages for children (< 18 , 312 samples), young adults ($[18 - 35)$, 1453 samples), middle aged ($[35 - 65)$, 1284 samples), and senior (> 65 , 286 samples).

Color FERET. This dataset [15] contains a total of 11K images taken from a semi-controlled environment, spanning 994 identities, and annotated with gender and race information. We filtered the images containing hard-profile poses (close to 90° in either direction) and processed them using the StyleGAN2 FFHQ face cropper to a 1024×1024 resolution. The final set used in this paper contains 994 identities and has 7.95K images, averaging 8 images per subject, only used for training PCA and LDA.

4.2. Latent editing and feature selection

We add a small scalar s steps along the age direction in W (Eq. 1), regenerate each image, and estimate its apparent age with MobileNetV2. Using the UTKFace test set split into 9 age groups, we compute the mean estimated age per group for each s , obtaining a discrete mapping $s \rightarrow \text{age}$. We then fit linear and polynomial curves per group to get a continuous relationship, which lets us pick the scalar s needed to reach any desired target age (older or younger).

We use two pre-trained face recognition models scoring over 95% accuracy in cross-age benchmarks, IResNet50 [17] and EdgeFace-S [14], to estimate identity preservation and MobileNetV2 age estimator model [26] to evaluate apparent age of the face images. We find verification threshold for both face recognition systems using the standard benchmark IJB-C [31] dataset. We use face recognition and age estimator to evaluate the effect of using linear features (from Linear SVR) compared to feature selection approaches aimed to preserve identity. We use UTKFace testing partition to compute the number of correctly verified age-manipulated samples and their age difference using the estimated age against the sample's origi-

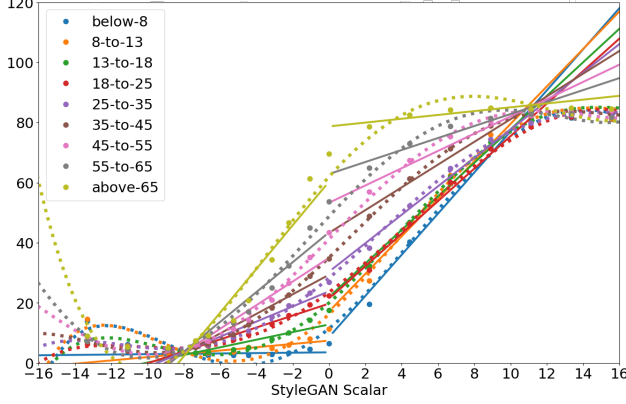


Figure 4. Baseline with linear and polynomial fitting for apparent age change across different age groups compared to the StyleGAN2 scalar change, using UTKFace test set. The scalar 0 denotes the original re-projected image. We note the differences in impact depending on the subject’s age group.

nal age label. With this difference, we calculate an *age gain* range, using its mean and one standard deviation, representing our confidence interval. We calculate the age gain for the aging/de-aging cases, depending on if we move the latent using a positive scalar to age the resulting face image or in the negative direction to de-age it. In the feature selection approaches, we set α from Eq. 3 to 1. We varied β from 0 (no change) to 0.25, 0.50, 0.75, and 1 to find the value that leads to the largest age gain preserving the identity. To that end, we compute the *age gain* for the four age group labels as described in Section 4.1. We interpolate the *age gain* at each rate of correctly verified samples from the test set of UTKFace, allowing us to examine the *age gain* range behavior at different face verification percentages.

5. Results

Using real faces with age labels from UTKFace [37], we apply StyleGAN2 latent space editing approach (see Section 3.1 to compute the parameters for identity preserving aging and de-aging. Then, we compare different feature selection strategies presented in Section 3.2. Finally, we generate fully synthetic dataset of identities using Langevin sampling [13] and aged/de-aged with our approach.

5.1. Apparent age estimation

Figure 4 illustrates the change in the apparent age for different scalar s from Eq. 1 computed separately for each of the 9 age groups (see Section 4.2 for details). The scalar value 0 corresponds to the re-projected original image, the negative scalar values correspond to de-aged images, and the positive scalar values correspond to the aged images. Figure 4 shows that the re-projection of the original faces (scalar value 0) appear younger as per the MobileNetV2 age

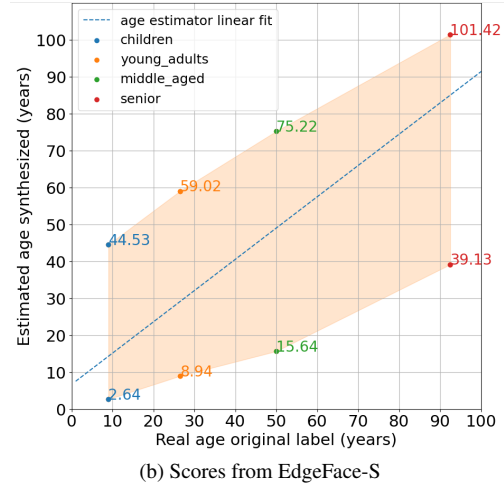
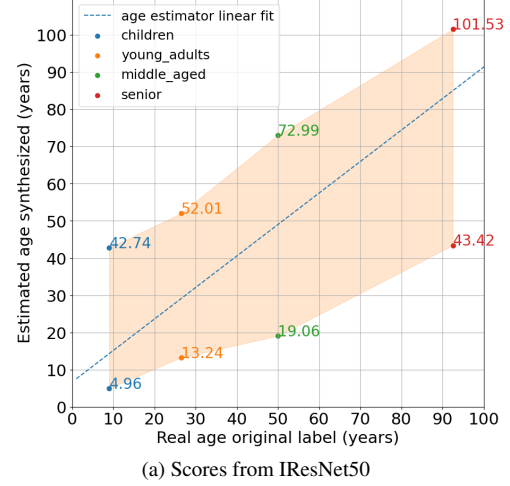


Figure 5. Range of estimated age from synthesized images, with subjects from distinct age groups while preserving identity on 75% of samples in UTKFace test set. We note the extended range of values in years with EdgeFace-S when compared to IResNet50, in both aging and de-aging cases.

estimator [26]. We notice non-linear behaviors when the face is aged or de-aged, especially, if the original age of the subject in the older or younger age groups. The figure also shows polynomial curves fitted to the data, which we can use to compute more accurate step-by-step increments when manipulating the age in the latent space.

5.2. Maximum aging ranges while preserving ID

We age/de-age faces with the baseline method and check identity with two face recognition models at cutoff rate of 75% of samples. We then linearly fit the estimated age of the original samples and hence find the age-gain for each age group. This age-gain is shown in Figure 5a as the highlighted region, within which aged/de-aged faces are matched with the original images using IResNet50 [17], and in Figure 5b using EdgeFace-S [14].

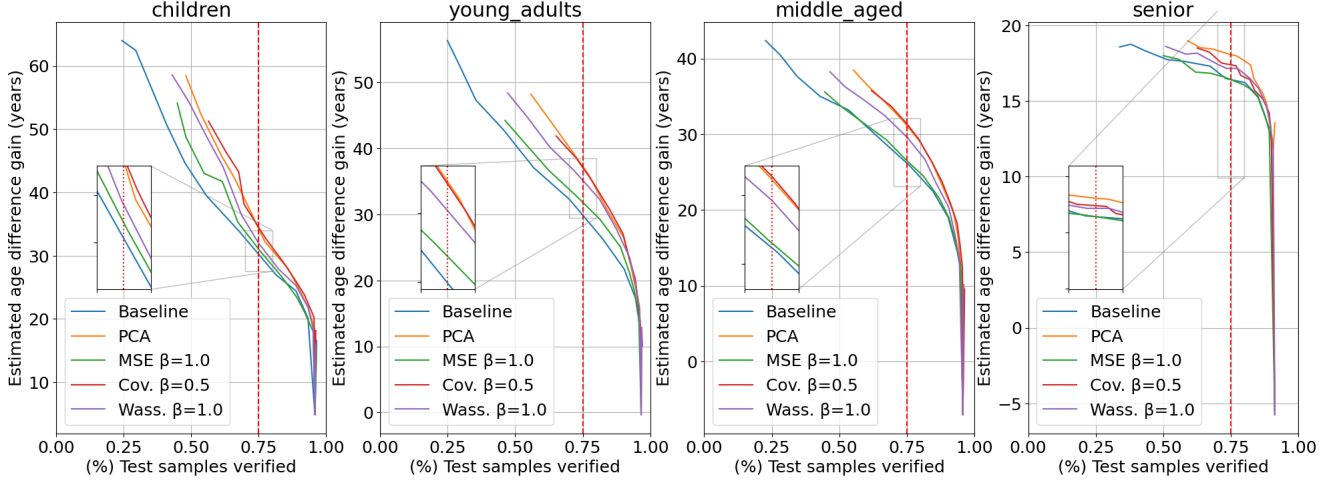


Figure 6. Age gain per number of samples verified when aging faces. We show face verification performance against age gain across different age categories. We highlight the differences at 75% of verified samples from the test set while using EdgeFace-S as face recognition system. The metrics for MSE, Covariance, and Wasserstein distance are computed using LDA bases for projection and re-construction.

Figure 5a of IResNet50 shows an obvious trend that the linear baseline approach can age children and de-age seniors by many years, while it does not work well in reverse. Also, middle age categories can be both aged and de-aged by a large margin. Figure 5b for EdgeFace-S shows that the apparent age range limits are wider across all categories. For example, EdgeFace can recognize children faces de-aged to 2 years and aged to 30 years, showing the impact of the chosen face recognition system.

5.3. Estimated age gain on feature selection methods

Figure 6 shows the performance of linear baseline approach against different feature selection mechanisms as described in Section 3.2. The scores from EdgeFace-S are sued as reference. The figure shows an expected drop in the rate of correctly recognized samples as we move in

the StyleGAN2 space, with respect to its apparent age (estimated) and allow us to compare ID-preservation performance of different feature selection strategies. Zoomed boxes emphasize the performance at 75% of recognized samples. All the feature selection methods improve over the baseline in terms of apparent age and the rate of samples correctly verified. In particular, PCA and LDA using Covariance show the most improvements over the linear baseline, with the latter being slightly better. These two approaches perform similarly, with covariance performing slightly better over PCA in the children age group, and PCA being more robust in the senior age group. Figure 7 compares our method and SAM [2], showing better performance qualitatively and quantitatively across larger age gaps.

5.4. Fully synthetic aged dataset

We publicly release a 20k-ID synthetic dataset¹ with 10 age versions (Figure 1), generated using Langevin sampling [13] and our SVR-based latent edit approach from Section 3.1.

6. Conclusion

We showed that a simple SVR-based latent direction and feature selection allows identity-preserving age edits in StyleGAN2 without re-training generators. Polynomial scalar-age mappings provide practical target-age control, and PCA/LDA masks improve the range of ID-preserving age edits. Future work includes alternative age estimators and latent navigation strategies.

Acknowledgement

This work was funded by InnoSuisse 108.708 INT-ICT.











Input	45 y/o	55 y/o	65 y/o	75 y/o
Ours				
				
FR score	-0.22	-0.25	-0.40	-0.66
SAM				
				
FR score	-0.24	-0.23	-0.23	-1.06

Figure 7. Age synthesis to specific target ages of our approach and SAM [2]. The face recognition (FR) score is the cosine distance of each image to the input using EdgeFace-S embeddings.

References

- [1] M. H. Abdulbas and N. D. Al-Shakarchy. Face information forensics analysis based on facial aging: A survey. In *2022 International Conference on Intelligent Systems and Computer Vision (ISCV)*, pages 1–7, 2022. [3](#)
- [2] Y. Alaluf, O. Patashnik, and D. Cohen-Or. Only a matter of style: age transformation using a style-based regression model. *ACM Trans. Graph.*, 40(4), July 2021. [4](#), [8](#)
- [3] G. Antipov, M. Baccouche, and J.-L. Dugelay. Face aging with conditional generative adversarial networks. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 2089–2093, 2017. [1](#), [2](#), [4](#)
- [4] D. Atkale, M. Pawar, S. Deshpande, and D. Yadav. Multi-scale feature fusion model followed by residual network for generation of face aging and de-aging. *Signal, Image and Video Processing*, 16, 04 2022. [2](#), [3](#)
- [5] S. Banerjee, G. Mittal, A. Joshi, C. Hegde, and N. Memon. Identity-preserving aging of face images via latent diffusion models. In *2023 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10, 2023. [1](#), [4](#)
- [6] P. K. Chandaliya and N. Nain. AW-GAN: face aging and rejuvenation using attention with wavelet GAN. *Neural Comput. Appl.*, 35(3):2811–2825, Sept. 2022. [2](#), [4](#)
- [7] P. K. Chandaliya and N. Nain. ChildGAN: Face aging and rejuvenation to find missing children. *Pattern Recognition*, 129:108761, 2022. [4](#)
- [8] P. K. Chandaliya, A. Sinha, and N. Nain. Childface: Gender aware child face aging. In *2020 International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–5, 2020. [4](#)
- [9] L. Colbois, T. d. Freitas Pereira, and S. Marcel. On the use of automatically generated synthetic image datasets for benchmarking face recognition. In *2021 IEEE International Joint Conference on Biometrics (IJCB)*, page 1–8. IEEE Press, 2021. [4](#)
- [10] D. Deb, D. Aggarwal, and A. K. Jain. Identifying missing children: Face age-progression via deep feature aging. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 10540–10547, 2021. [4](#)
- [11] M. Falkenberg, A. B. Ottsen, M. Ibsen, and C. Rathgeb. Child face recognition at scale: Synthetic data generation and performance benchmark. *Frontiers in Signal Processing*, 2024. [4](#)
- [12] H. Fang, W. Deng, Y. Zhong, and J. Hu. Triple-GAN: Progressive face aging with triple translation loss. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 3500–3509, Los Alamitos, CA, USA, jun 2020. IEEE Computer Society. [3](#)
- [13] D. Geissbühler, H. O. Shahreza, and S. Marcel. Synthetic face datasets generation via latent space exploration from brownian identity diffusion. In *Proceedings of the 42nd International Conference on Machine Learning (to appear)*. PMLR, 2025. [1](#), [3](#), [7](#), [8](#)
- [14] A. George, C. Ecabert, H. O. Shahreza, K. Kotwal, and S. Marcel. Edgeface: Efficient face recognition model for edge devices. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2024. [6](#), [7](#)
- [15] M. M. Ghazi and H. K. Ekenel. A Comprehensive Analysis of Deep Learning Based Representation for Face Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 102–109, Los Alamitos, CA, USA, July 2016. IEEE Computer Society. [6](#)
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks. *Commun. ACM*, 63(11):139–144, Oct. 2020. [1](#)
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [6](#), [7](#)
- [18] G.-S. J. Hsu, R.-C. Xie, Z.-T. Chen, and Y.-H. Lin. Agetransgan for facial age transformation with rectified performance metrics. In *European Conference on Computer Vision*, 2022. [4](#)
- [19] Z. Huang, S. Chen, J. Zhang, and H. Shan. PFA-GAN: Progressive face aging with generative adversarial network. *IEEE Transactions on Information Forensics and Security*, 16:2031–2045, 2021. [3](#)
- [20] A. Kale and O. Altun. Face age synthesis: A review on datasets, methods, and open research areas. *Pattern Recognition*, 143:109791, 2023. [3](#)
- [21] T. Karras, S. Laine, and T. Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 43(12):4217–4228, Dec. 2021. [1](#)
- [22] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. Analyzing and Improving the Image Quality of StyleGAN. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8107–8116, Los Alamitos, CA, USA, June 2020. IEEE Computer Society. [2](#)
- [23] P. Korshunov, H. Chen, P. N. Garner, and S. Marcel. Vulnerability of automatic identity recognition to audio-visual deepfakes. In *2023 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10, 2023. [1](#)
- [24] P. Korshunov, A. George, G. Özbülak, and S. Marcel. Vulnerability of face age verification to replay attacks. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4830–4834, 2024. [1](#)
- [25] P. Korshunov, K. Kotwal, C. Ecabert, V. Vedit, A. Mohammedi, and S. Marcel. Investigation of accuracy and bias in face recognition trained with synthetic data. In *IEEE International Joint Conference on Biometrics (IJCB)*, Sept. 2025. [1](#)
- [26] P. Korshunov and S. Marcel. Face anthropometry aware audio-visual age verification. In *Proceedings of the 30th ACM International Conference on Multimedia, MM '22*, page 5944–5951, New York, NY, USA, 2022. Association for Computing Machinery. [1](#), [6](#), [7](#)
- [27] B. Li, S. Cai, W. Liu, P. Zhang, Q. He, M. Hua, and Z. Yi. Dystyle: Dynamic neural network for multi-attribute-conditioned style editings. In *2023 IEEE/CVF Winter Con-*

- ference on Applications of Computer Vision (WACV), pages 189–197, 2023. 4
- [28] L. S. Luevano, Y. Martínez-Díaz, H. Méndez-Vázquez, M. González-Mendoza, and D. Frey. Swiftfaceformer: An efficient and lightweight hybrid architecture for accurate face recognition applications. In A. Antonacopoulos, S. Chaudhuri, R. Chellappa, C.-L. Liu, S. Bhattacharya, and U. Pal, editors, *Pattern Recognition*, pages 244–258, Cham, 2025. Springer Nature Switzerland. 1
- [29] L. S. Luevano, Y. Martínez-Díaz, H. Méndez-Vázquez, M. González-Mendoza, and D. Frey. Assessing the performance of efficient face anti-spoofing detection against physical and digital presentation attacks. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1021–1028, 2024. 1
- [30] Y. Martínez-Díaz, H. Méndez-Vázquez, L. S. Luevano, and M. Gonzalez-Mendoza. Exploring the effectiveness of lightweight architectures for face anti-spoofing. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 6392–6402, 2023. 1
- [31] B. Maze, J. Adams, J. A. Duncan, N. Kalka, T. Miller, C. Otto, A. K. Jain, W. T. Niggel, J. Anderson, J. Cheney, and P. Grother. Iarpa janus benchmark - c: Face dataset and protocol. In *2018 International Conference on Biometrics (ICB)*, pages 158–165, 2018. 6
- [32] X. Ning, S. Xu, F. Nan, Q. Zeng, C. Wang, W. Cai, W. Li, and Y. Jiang. Face editing based on facial recognition features. *IEEE Transactions on Cognitive and Developmental Systems*, 15(2):774–783, 2023. 4
- [33] S. Palsson, E. Agustsson, R. Timofte, and L. Van Gool. Generative adversarial style transfer networks for face aging. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2165–21658, 2018. 3
- [34] X. Tang, Z. Wang, W. Luo, and S. Gao. Face aging with identity-preserved conditional generative adversarial networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7939–7947, 2018. 2, 4
- [35] W. Wang, Z. Cui, Y. Yan, J. Feng, S. Yan, X. Shu, and N. Sebe. Recurrent face aging. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2378–2386, 2016. 3
- [36] Z. Wang, Z. Liu, J. Huang, S. Lian, and Y. Lin. How old are you? face age translation with identity preservation using gans. *arXiv preprint arXiv:1909.04988*, 2019. 3
- [37] Z. Zhang, Y. Song, and H. Qi. Age progression/regression by conditional adversarial autoencoder. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4352–4360, Los Alamitos, CA, USA, jul 2017. IEEE Computer Society. 6, 7
- [38] S. Zhou, K. Chan, C. Li, and C. C. Loy. Towards robust blind face restoration with codebook lookup transformer. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 30599–30611. Curran Associates, Inc., 2022. 6
- [39] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251, 2017. 3