



A DEEP DIVE INTO STYLE-BASED AGE MANIPULATION: REVISITING SAM FOR LIFESPAN FACE EDITING

Swathi Gowroju¹, Devineni Gireesh Kumar^{*2}, S A Md Noorulla Baig³,
K. Moulya Reddy⁴, P. Devika⁵ and D. Sai Chand⁶

^{1, 2}Associate Professor, Department of CSE(AI&ML), Sreyas Institute of Engineering and Technology, Hyderabad, 500068, Telangana, India.

³Associate Professor, Department of CSE (AI), Vemu Institute Of Technology, P Kothakota, Chittoor - 517 112, AP, India.

^{4, 5, 6}UG Student, Department of CSE(AI&ML), Sreyas Institute of Engineering and Technology, Hyderabad, 500068, Telangana, India.

¹<https://orcid.org/0000-0002-4940-1062>, ²<https://orcid.org/0000-0003-0575-471X>, ³<https://orcid.org/0009-0004-6779-042X>

⁴<https://orcid.org/0009-0001-2243-7698>, ⁵<https://orcid.org/0009-0001-1257-6598>, ⁶<https://orcid.org/0009-0008-0751-9100>

Email: swathigowroju@gmail.com, *gireesh218@gmail.com, noorulla.baig@gmail.com, moulyareddy370@gmail.com,
palsamdevika@gmail.com, saichandrathod003@gmail.com

ARTICLE INFO

Article History

Received: February 14, 2026

Reviewed: April 23, 2026

Accepted: May 23, 2026

Published: June 30, 2026

Keywords:

Facial Age Transformation,
StyleGAN2,
Identity Preservation,
Regression,
Latent Space

ABSTRACT

Change of facial age aims to realistically alter a person's photograph's age without loss of identity, a very important aspect for forensic work, entertainment, and all digital media. The Style-Based Age Manipulation (SAM) model is based on the expressive latent space of StyleGAN2 to enable photorealistic, seamless age progression and regression across the human life cycle (5 to 100 years). SAM uses a pSp encoder to embed real images into the W + latent space and a new ageing encoder for age-conditioned, non-linear transformations with a fixed StyleGAN2 generator for high-quality synthesis. LL_2 , LPIPS, ArcFace, age regression, and cycle consistency losses have been optimized in SAM to improve visual quality, identity fidelity, and reversible transformations. Its novelty lies in a regression-based architecture that enables continuous, disentangled age control, in contrast to previous methods that rely on discrete age bins or linear latent traversals. SAM outperforms on LPIPS (mean: 0.11), ArcFace similarity (mean: 0.93), and age classification accuracy ($91\% \pm 5$ years). Qualitative results provide lifelike smooth age transitions, while quantitative measures show 10% improved identity preservation and 18% better visual fidelity compared with CNN baselines. A strong, generalizable method is developed in SAM and sets a new standard for identity-holding facial age change, with broad practical applicability.



Copyright ©2026 by authors and Galileo Institute of Technology and Education of the Amazon (ITEGAM). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

I. INTRODUCTION

Face ageing is a fascinating computer vision challenge, with promising applications in entertainment, forensics, digital media, and identity verification [1], [2]. The objective is to alter the perceived age of a subject in an image without affecting their identity and maintaining visual quality [3], [4]. Traditional methods for face ageing rely on narrow age intervals and fail to preserve fine facial details or to generate high-resolution images [5], [6]. With the advent of deep generative models, particularly Generative Adversarial Networks (GANs) [7], face ageing algorithms are now of much better quality and versatility [8], [9]. StyleGAN, a state-of-the-art GAN architecture with high-resolution image synthesis and a dense latent space, enables subtle face manipulation [10]. However, current methods built on StyleGAN generally either require manual latent traversal or only support binary-age transitions (e.g., from child to adult). Also, they fail to fully disentangle the ageing process from facial expressions, poses, and hairstyles. To address these concerns, this work aims at Style-Based Age Manipulation (SAM). This recently introduced regression-based approach converts facial images and target age values into the StyleGAN latent space [11]. SAM-linear smooth, continuous, and controllable ageing by learning a non-linear, disentangled path guided by an age regressor [12]. This results in photorealistic ageing effects across the full human lifespan without compromising identity or unwanted changes.

I. 1 MOTIVATION

Face ageing is socially conspicuous and visually rich in computer vision and artificial intelligence [13]. As surveillance systems and social media increasingly provide large volumes of facial images, there has never been a greater demand for realistic and reliable age transformation [12]. Traditional deep learning methods, e.g., CNN-based or fixed-domain translation GAN-based [14], work by projecting facial attributes onto discrete bins of age (e.g., 10s, 30s), which is restrictive in terms of expressiveness and generalisation [5]. Moreover, these models will transform unwanted aspects such as facial expressions, hair colour, or even gender, which compromises identity consistency [3]. The motivation for SAM stems from the flaw in such earlier approaches. Unlike methods that rely on categorical regression or latent vector interpolation, SAM treats age as a continuous regression problem and therefore accommodates smooth, natural progression or regression of age [11]. It leverages the highly disentangled, semantically rich latent space of StyleGAN2, where fine-grained features such as wrinkles, facial shape, and skin texture can be manipulated without altering the underlying identity [10]. Furthermore, SAM eliminates the need to manually construct traversal directions because it learns an explicit mapping from the input image and the target age to the corresponding latent code [15]. This design not only improves visual realism but also enhances identity preservation [16], which is critical for applications such as forensic age progression, AI-based age verification systems, and digital entertainment [6].

I. 2 OBJECTIVE 8AND SCOPE

The purpose of this study is to compare other works and demonstrate the capabilities of the Style-Based Age Manipulation (SAM) model for realistic aged faces in high resolution [9]. The directions of the study are as follows: Investigating SAM Architecture: Facing the issues of how SAM combines an age encoder with a StyleGAN2 generator and an auxiliary network (ArcFace and age regressor) [16], [10] for producing fine-tuned identity-preserving results. Continuous Age Transformation: Contrary to usual models that use fixed age ranges [14], SAM creates faces spanning 5 to 100 years, enabling smooth, natural transitions [11]. Could the present work address how well SAM does this and its ability to generalise across different kinds of faces? Preservation of Non-Age Features: Another stream of the study is testing SAM for preserving non-age features such as pose, hair, and expression from being altered, thereby enhancing realism or avoiding uncanny effects [17]. Real-World Applications: By presenting sample results, the paper will also discuss real-world applications of SAM, including digital character creation, age progression for missing people, and secure identity management systems [6]. SAM aims to generalize and fill the gaps between both directions [11].

I. 3 NOVELTY

Through Style-Based Age Manipulation, the idea becomes novel, as it reframes the task of facial age manipulation into a continuous regression problem in the expressive latent space of StyleGAN2. In contrast to earlier discrete solutions for age categories, such as StarGAN or CUSP, which restrict transformations to usually-imposed bins (like 10s, 20s), the SAM allows for fine control over the entire human age spectrum, from 5-100 years of age, with the possibility of smoothly and naturally transitioning to any target birth age, even fractional values. Mimicking Age Flow in continuous age control, the proposed regression-based scheme diverges from classical approaches. It uses non-linear transformations in the $W+$ latent space, avoiding the sudden changes observed in domain-translation models. SAM's true novelty hinges upon its ageing encoder, based on the principles of residual learning, that predicts age-specific residual latent vectors that would only localize changes such as wrinkles or skin texture while being agnostic to identity-critical features such as eye shape or nose structure, which is a challenge that linear latent traversal methods like InterFaceGAN struggle to overcome due to attribute entanglement. Combining a well-trained pSp encoder to perform robust image-to-latent mapping with a frozen generator from StyleGAN2 for high-fidelity synthesis enables computational efficiency, whereby SAM trains only the ageing encoder and remains competitive with state-of-the-art performance.

II. RELATED WORK

Over the last two decades, the transformation of facial ageing has undergone significant changes, from traditional image-processing-based methods to deep-learning methods [1], [8]. Recent developments in GANs, especially StyleGAN, have greatly improved the quality, realism, and control over synthesized facial images [10]. This section reviews existing works in facial age synthesis, which can be categorized into image-to-image translation models and latent space manipulation approaches, both of which lie at the heart of the approach proposed in SAM [17], [18]. This face-to-face ageing problem has attracted researchers' interest for many years due to its complexity and broad applicability [13]. The field has changed as it has grown from rule-based to data-driven deep learning techniques [7]. This section discusses a couple of important contributions and categorizes the literature into two broad categories: (1) image-to-image translation-based models and (2) GAN latent space. A detailed survey of the literature is provided in Table 1.

Table 1: Survey of literature on style-based age manipulation.

[1]	2024	DLAT ⁺	Proposed a life span-reaching facial age transformation model along with the IDAG to augment for identity variation.	Directly influenced SAM's identity preservation and age advancement through its regression build-up.
[3]	2024	Synthetic Face Ageing	Studied the extent to which synthetic ageing impacts recognition in the face domain, keeping minimum identity drift an important issue in efficient recognition across the ages.	Emphasised identity preservation but influences SAM usage of ArcFace loss for this effect.

[2]	2024	3D Uplifting	Wanted to achieve an approach for facial age transition inspired by 3D and uplifting features at different stages of life.	Influenced SAM's realism emphasis, yet its 3D nature led us to edit with 2D StyleGAN2 latent refinement.
[4]	2024	ISA-GAN	Designed an inception-based self-attentive encoder-decoder GAN for face-to-face generation by describing face photographs.	Inspired by SAM's encoder architecture, although modified for continuous control on age in latent space.
[8]	2024	AgeSynthGAN	Realistic facial age-changing about different age shifts.	Shaped SAM's StyleGAN2 base, though we extended this into a continuous regression approach for flexibility.
[13]	2024	Ceph2Style	Prediction in facial growth direction is based on editing the StyleGAN latent space, with emphasis on consistency with respect to time.	Motivated latent space editing of SAM, though we placed our attention on bidirectional age mapping.
[5]	2024	Face Re-Aging	Suggested temporally consistent face-ageing, solving identity preservation over time.	SAM, therefore, influenced temporal consistency aspects but was further improved by a regression running way.
[6]	2024	De-Ageing Networks	Devised de-ageing networks for forensic sketching, involving precise age reversal while preserving his identity.	Extended inspiration to age at SAM in both directions, but with greater emphasis on lifespan coverage than on forensic specifics.
[12]	2025	Text-to-Image Models	Evaluating the text-to-image generative models to preserve age representation truthfully, but they are limited in controlling ages.	Guided SAM's image-based age regression while not complicating it with text.
[9]	2024	Re-Ageing GAN	Engineered a personalized age change GAN, targeting identity conservation as opposed to age groups.	Influenced SAM's architectural design to be personalized and identity- preserving, optimized using StyleGAN2.

Source: Authors, (2026).

Image2StyleGAN++, introduced by [15], grants the capability to move real images to the latent space of StyleGAN, thereby enabling high-fidelity editing; the focus of Image2StyleGAN++ rested on the correct reconstruction and manipulation of real-face images while retaining fine details, all contingent on operation within the latent space. This work was vital to transferring real images into the W^+ latent space using the pSp encoder in the SAM domain while allowing variation in the input flow. However, SAM introduced an additional dimension by incorporating age-related changes within an augmented regression framework for continuous control. ReStyle is an encoder for the residual style of a StyleGAN developed by [19] and applying a refinement iteration to achieve accurate latent codes. The fidelity of the image-to-latent mapping has improved with an iteratively updated latent representation. ReStyle opened the door for SAM to investigate residual learning, particularly regarding the design of the ageing encoder, to refine latent codes conditioned on the target age. SAM ultimately derived its inspiration from the notion of continuous age regression for real-time processing and efficiency, rather than iteration. The researchers [11] designed AgeFlow, an architecture for conditional age progression and regression with continuous control for age conditioning using normalizing flows. AgeFlow enables fine-grained control over age by modelling age transformations as smooth probabilistic processes. The research inspired SAM's regression-based framework for continuous age manipulation across the lifespan, defined as 5-100 years. Instead of using the flow-based approach employed by AgeFlow, SAM uses the latent space of StyleGAN2 to achieve photorealistic outputs while preserving identity.

By [14] described StarGAN v2 as a unified model of multi-domain image translation applicable to transformations in specific attributes, including age. This was modelled to transform images with a single generator into many domains. Using the StarGAN v2 model, SAM was built for continuous age control. Unlike the star GAN v2 approach of discrete domain transformations, SAM therefore used regression analysis to effect smoother, finer transitions. A study by [17] discussed local semantics in the latent space of StyleGAN, enabling the independent editing of specific facial attributes. Disentangled control for accurate manipulation tasks, such as age manipulation, was emphasized in this work, which, in turn, motivated SAM to focus on disentangling age from other facial attributes, devoid of any influence on identity or background transformations, such as wrinkles or skin texture. Additionally, SAM implemented a non-linear regression framework for lifespan coverage. A very important loss function, ArcFace, was introduced by [16] in conjunction with the radius margin, which has been shown to mitigate identity loss in facial images. Given that an important factor in preserving identity throughout transformations is what made ArcFace most suited for the application in SAM, where it is incorporated as a loss function to endorse conservation of facial identity during the processes of age progression versus regression, directly confronting the identity drift problems of previous models and broadening the relevant domain of applicability of SAM.

GANSpace was introduced by [20] as a method for discovering interpretable latent directions in GANs that enable control of semantic attributes such as age or expression. The intuitive attribute editing enabled by GANSpace stems from its analysis of the latent space. The findings of this work influenced the manipulations of SAM's latent space, and therefore, the design of the ageing encoder was set to navigate the W^+ space for age-specific edits. However, SAM executed fine control over continuous transformations of age via a nonlinear regression approach. By [10] introduced StyleGAN2 as an improved version of StyleGAN with better image quality and disentangled manipulations over the W^+ latent space. Its architecture allows for high-fidelity synthesis of facial images and vigorous editing in the latent space. It is the foundation generator for SAM, providing an expressive latent space for photorealistic age transformations. SAM then incorporated a regression-ageing encoder to optimize continuous age control. By [21] developed StyleCLIP, a text-guided manipulation framework that enables natural-language control of StyleGAN prisms. This was one example of conditioning latent space edits. Although it's a bit revealing within StyleCLIP, whether SAM employed it to disclose the age-conditioning mechanism, SAM is much more precise, with crisp, continuous regressions that justify itself via its regression-based model and inputs. By [22] even described the pSp encoder, which acts as a pixel-to-StyleGAN encoder, mapping real images into the W^+ latent space of StyleGAN for image-to-image translation. When it comes to how the pSp encoder could accommodate realistic images, this bit was what made SAM's design for the first part of generating latent code for real images. Further, it incorporated the ageing encoder into the final process to ensure such codes undergo continuous transformations over age.

By [18] introduced and explained InterFaceGAN, which aims to identify linear paths in the GAN latent space to enable semantic attribute editing, for example, age. The method showed that one could reasonably control the manipulation of an attribute but assumed linearity. Thus, the non-linear latent transformation approach articulated by SAM was partly inspired by InterFaceGAN and now supports continuous age regression to address the limitations of linear traversals and to handle complex ageing effects. According to [23] propose the framework StyleRig, which enables rigging latent-space edits for targeted manipulation in the 3D action context, yielding 3D control over the StyleGAN-generated portraits. Though for 2D age manipulation in the StyleGAN2 latent space, StyleRig paved the way for SAM's latent editing in the parametric space of facial attributes. By [24] presented an encoder tailored for fine-grained image manipulations based on StyleGANs that optimizes the generation of latent codes for attribute editing. This made defining the manipulation targeted as the relevant consideration in designing the encoder architecture. The influence trickled down to SAM's ageing encoder, which conditions the latent codes on target age for continuous regression. SAM, built to be real-time and cover a lifespan, was pushed to achieve practical optimization. StyleSpace by [25] provides a disentanglement analysis of controls in the latent space of StyleGAN for manipulation of very fine-grained attributes. Specific channels were revealed for attributes such as age, allowing them to be edited. This research helped guide SAM in disentangling age from other facial features so that transformations such as wrinkles or hair color would not contaminate identity or background, which are important for producing realistic outputs. CycleGAN, the cycle-consistent adversarial network for unpaired image-to-image translation, was designed by [7]. It allows transformations without paired data and includes a cycle-consistency mechanism to enforce the reversibility of transformations. CycleGAN also inspired SAM's development of a cycle consistency loss, which enforces non-drift for structure-specific transformations during age progression and regression.

II.1 IMAGE-TO-IMAGE TRANSLATION MODELS FOR AGE PROGRESSION

Early attempts at face ageing were mostly template-driven, based on 3D morphable models or statistical patterns of craniofacial ageing. These methods were not personalized and generated general ageing outcomes irrespective of one's own features. The introduction of GANs marked a turning point in the field. CycleGAN, designed for unpaired image-to-image translation, enabled researchers to translate age domains (e.g., "young" to "old") without paired data. C-AAE (Conditional Adversarial Autoencoder) and AgecGAN pursued the same strategies, learning to encode inputs into several age classes using conditional inputs. StarGAN took a step ahead by proposing an end-to-end multi-domain translation model, with a single generator that can convert between any two of numerous possible domains (e.g., age 10, age 20, age 30, age 40, etc.). These models, however, were not without their downsides: No continuous age representation: They could only support a few discrete classes, leaving no room for fine-grained control. Overfitting to domain labels: Transitions were not smooth and had no subtle variation. Identity drift: Unnecessary changes in facial structure would most probably result in identity loss. Lifespan-Aware Face Ageing (LIFE) and CUSP solved these problems by creating more accurate age groups and lifespan-aware loss functions. However, they still lacked continuous control and required dense age-labelled training data. On the other hand, SAM recasts face ageing as a regression problem rather than a classification problem, thereby avoiding the rigidity of domain-based methods. It enables transformation to any desired target age (for instance, age 63.5 rather than a specific group), enhancing realism and user control.

II.2 LATENT SPACE MANIPULATION AND GAN INVERSION

Another category of face ageing is latent space manipulation, where real face images are embedded in the latent space of a GAN, predominantly StyleGAN, and attribute editing is performed by manipulating latent vectors. InterFaceGAN proposed a method for discovering linear directions in latent space that corresponds to semantic attributes such as age, gender, or expression. While intuitive, these approaches often assume a linear relationship between features and latent directions, which is not always true for complex transformations such as ageing. Struggle with attribute entanglement, where changing one aspect changes others inadvertently (e.g., ageing might change the background or hairstyle). Require manual tuning or traversal to achieve optimal results. Rely on GAN inversion, which is prone to reconstruction artefacts and limit's identity consistency. Latent-Transformer and StyleFlow were trying to enhance editing accuracy by learning mappings from target attributes to latent codes. However, they were still interpolating between latent vectors and had poor inversion quality, especially for real, high-resolution face inputs. SAM addresses these limitations with a strict end-to-end solution: An encoder (pSp-based) embeds real images into the StyleGAN2 W^+ latent space. An age encoder learns a conditioned residual transformation in latent space, conditioned on the target age. LPIPS and identity losses (ArcFace) are employed to compel visual quality and identity preservation.

The regression network-informed age loss guarantees accurate alignment at the target age without the inaccuracy inherent in traversal-based methods in the latent space. This tight integration of GAN inversion, latent code manipulation, and explicit age conditioning makes SAM one of the strongest and most realistic face-ageing models to date.

III. METHODOLOGY

The Style-Based Age Manipulation (SAM) model is designed to perform continuous, realistic facial ageing and regression without regard for the subject's identity or other key facial features. Working entirely within the StyleGAN2 latent space, it uses age guidance while employing a novel residual learning mechanism to preserve identity. Essentially, age transformation in the face domain by SAM leverages the powerful latent space of StyleGAN2. Whereas classical methods operate within a linear trajectory in the latent space or perform domain translation across age groups, SAM prescribes a regression-based, end-to-end approach for continuous, identity-preserving age transformation. This section provides in-depth explanations of SAM's architecture, training paradigm, and loss functions.

III.1 OVERVIEW OF THE SAM FRAMEWORK

SAM tackles age progression as a conditional image-to-image translation problem. It is composed of an ageing encoder that takes the input image and the desired age as input, then outputs a chain of style vectors in the $W+$ latent space of StyleGAN2. These vectors are fed into a frozen StyleGAN2 generator to render the age-transformed image. For model training, an encoder (pSp) is used to provide a base latent code from the input image, thereby preserving reconstruction fidelity. The SAM encoder learns a residual transformation to fine-tune the base latent code to the desired age shift during training.




Figure 1: SAM Architecture.

Source: Authors, (2026).

Figure 1 shows the architecture and training pipeline of the Style-Based Age Manipulation (SAM) model. The process starts with an input image x and a target age α_t . The image is fed through two independent encoders: A frozen identity encoder produces the identity-preserving latent code w^* , where an ageing encoder learns to generate a residual latent vector Δw based on the target age. These vectors are used to generate the final latent vector w^+ , which is passed to the fixed StyleGAN2 generator to generate an aged output image y_{out} . To realize high-quality and precise age transformation, four loss functions are adopted: Age Loss to match with the target age, Identity Loss to maintain facial identity, Perceptual Loss for realism of the image, and Cycle Consistency Loss to provide reversibility of the transformation.




Figure 2: Facial Age Transformation Process in SAM.

Source: Authors, (2026).

Figure 2 shows the important steps of the Style-based Age Manipulation (SAM) method, which provides high-resolution identity-preserving face ageing. The method uses the latent space of StyleGAN2 and employs independent encoders for age conversion and identity preservation. The procedure starts by concatenating the input face image and the target age condition, which can be either a scalar or a map (i.e., an age embedding). This concatenation is used as input to the ageing network, which is trained to adjust age-related features within the image. This conditioning allows the model to generate outputs that correspond to certain age targets. The age encoder takes the concatenated input and encodes an age-specific latent residual vector (Δw). It is the semantic transformation needed in StyleGAN's latent space to capture the target age, e.g., wrinkles, skin texture, or facial contours. Importantly, this residual is learned in a way that avoids interference with other properties, such as identity, pose, and lighting. Simultaneously, the face image input is fed through a frozen identity encoder, which embeds it in StyleGAN's latent space (W) to get a latent vector w^* . The code retains identity-preserving details so that the subject's appearance ages, while its underlying identity remains unchanged. The residual age Δw is added to the underlying latent code w^* to produce a resulting latent vector $w^+ = w^* + \Delta w$. The combined vector simultaneously represents both target age and retained identity. This operation is significant because it decouples age from identity while still maintaining the full expressiveness of StyleGAN's latent space. The resulting latent vector w^+ is subsequently fed into a fixed StyleGAN2 generator, which transforms it into a high-resolution aged face image. Since StyleGAN2 generates photorealistic facial images, the resulting output retains fine-grained details, making the aged version look natural, sharp, and coherent.

Table 2: Description of Modules in the SAM Model Pipeline.

Module	Function
Input Image x	Original facial image provided as input
Target Age at	Desired target age to guide the ageing transformation
Frozen Identity Encoder	Extracts identity vector w^* , remains unchanged during training.
Ageing Encoder E_{age}	Learns age-specific residual changes in latent space
Latent Space Residual Δw	Represents the age transformation component in the latent space
Final Latent Vector w^+	Sum of identity and age vectors used for image synthesis
Fixed StyleGAN2 Generator	Generates a high-resolution aged image using w^+
Aged Face Output y_{out}	Output aged image with applied age transformation
Loss Functions	Combines Age Loss, Identity Loss, Perceptual Loss, and Cycle Consistency Loss to train the model effectively

Source: Authors, (2026).

As evidenced by Table 2, the SAM model comprises a series of interdependent modules intended to unpick identity and age characteristics and blend a photorealistic aged face. The role of each module in the training pipeline is unique: it contributes to the production of age-progressed faces while preserving the subject's original identity. The target age and the input image are processed simultaneously: one uses a frozen identity encoder to maintain identity, and the other uses an ageing encoder to analyse the changes in the latent space corresponding to ageing. The calculated latent space residual (Δw) is indicative of the variation that occurs per age and which, when added to the identity-maintaining vector (w^*), creates a final latent code (w^+). The final vector is produced by a StyleGAN2 generator and used to generate a high-resolution aged face. To render realistic and identity-consistent output, the model is trained by utilising several loss functions: age loss, identity loss, perceptual loss, and cycle consistency loss. This multi-layered architecture, as shown in the figure and table, demonstrates SAM's capacity to learn fine-grained age changes while reducing distortion in other facial features, such as pose and expression.

III.2 AGING ENCODER

An ageing encoder can predict an insert to transform the input face to the target age. It comprises a Feature Pyramid Network (FPN) architecture that generates 18 independent 512-dimensional vectors, equal in number to the style layers in StyleGAN2's W^+ space. The encoder takes as input a 4-channel image, namely an RGB image concatenated with the target age broadcast as a constant-valued age map, and outputs a residual vector that is added to the base latent code w^* generated by the pSp encoder. This modified latent code is then passed to StyleGAN2 to synthesise an image. Since the residual vector is used to incorporate only age-related changes into the latent code, while identity and image structure are preserved, this structure allows the ageing encoder to perform only age-related transformations.

III.3 LOSS FUNCTIONS

SAM uses a blend of multiple loss functions to enforce age accuracy, identity consistency, and visual realism. L_2 Loss: Facilitates pixel-wise similarity between reconstruction and input in the cycle consistency step. LPIPS (Perceptual Loss): Fosters perceptual similarity by using deep features from a VGG network. It avoids unnatural shifts and fuzzy outputs. Identity Loss (L_{id}): Applies an ArcFace model to compute the cosine similarity between the identity embeddings of the output and input faces. The age gap scales it—greater age gaps permit greater identity drift.

$$L_{id} = 1 - \cos(f(x), f(x'))$$

Age Loss (L_{age}): aligns the output face with the target output age using an age regressor. L_2 distance between the predicted and target age anchors the regression performance.

$$L_{age} = (R(x') - a_{target})^2$$

Regularisation Loss (L_{reg}): Maintains the latent code near the mean latent vector to prevent drifting away from the latent manifold, causing artefacts.

$$L_{reg} = \sum_{i=1}^d (w_{final,i} - w'_i)^2$$

Cycle Consistency Loss (L_{cycle}): The model undergoes a round-trip: from an image to a target age, then back to the original age. The reconstructed image is compared with the original to minimise loss of identity and structure.

$$L_{cycle} = \sum_{i=1}^n |x_i - x'_{cycle,i}|$$

The loss function overall is a weighted sum of all the components:

$$L_{total} = \lambda_{L2} \cdot L_{L2} + \lambda_{LPIPS} \cdot L_{LPIPS} + \lambda_{reg} \cdot L_{reg} + \lambda_{id} \cdot L_{id} + \lambda_{age} \cdot L_{age} + \lambda_{cycle} \cdot L_{cycle}$$

III.4 TRAINING STRATEGY

Building Blocks (Frozen): StyleGAN2 generator. ArcFace identity model. VGG-based LPIPS model. Age regression network—pSp encoder.

Trainable: The Ageing Encoder (E_{age}) is the only trainable building block.

Training Process: Input image x is encoded to latent vector w^* with the pSp encoder. Target age at is sampled at random from [5, 100] and paired with x . The ageing encoder predicts a residual to be added to w^* . Finally, latent code is passed through the generator to generate y_{out} . A cycle consistency pass is done by converting y_{out} back to the initial age as . The model is trained with the Adam optimiser and learning rate scheduling. Training is performed on high-resolution face datasets (e.g., CelebA-HQ) at 1024×1024 resolution.

IV. RESULTS AND DISCUSSIONS

This section provides an extensive assessment of the Style-Based Age Manipulation (SAM) model, a new deep learning model designed for continuous, identity-consistent facial age transformation across the human lifespan (5–100 years). Based on the theoretical groundwork of StyleGAN2's expressive latent space and regression-based method of age manipulation, SAM combines a frozen pSp encoder, trainable ageing encoder, and StyleGAN2 generator under the direction of a set of loss functions (L_{L2} , LPIPS, ArcFace, age regression, and cycle consistency) to attain photorealistic ageing results. The assessment combines qualitative visualisations, for example, age progression timelines, with quantitative measures, for example, pixel-wise. Experiments were performed on a single input image to produce aged outputs at target ages (0, 10, ..., 100), which were validated using a ResNet-34 age classifier and compared with a CNN baseline. The results highlight SAM's advantage in smooth age transitions, identity maintenance, and high-resolution output generation, validating its use in digital forensics, entertainment, and identity verification, as discussed in Section 3.

IV.1 THEORETICAL FOUNDATION AND EXPERIMENTAL SETUP

The theoretical background for the operation of SAM is in its capacity to represent facial ageing as continuous regression in the disentangled W^+ latent space of StyleGAN2, in contrast with earlier discrete regression models that were predominant. The mapping is established through a non-linear regression learned by the ageing encoder that supersedes a base latent code (w^*) from a pSp encoder to encode target age features (e.g., wrinkles, skin texture) while maintaining identity through ArcFace loss and perceptual realism through the LPIPS loss. The generator of StyleGAN2 is fixed for image generation, and the cycle consistency loss helps reverse the process, thereby helping preserve against structural drift. Unlike the prevalent CNN-based architecture that relies heavily on convolutions and linear feature mappings, our design enables the encoding of more complex ageing effects while preserving identity. The experimental procedure employed a single RGB facial image (256x256 pixels) downloaded to a local directory, aligned with the dlib 68-landmark predictor to evenly distribute features across the image. The SAM model generated aged output at eleven target ages (0, 10, ..., 100). Resizing to 256x256, normalisation to the pixel value range [-0.5, 0.5], and age-wise transformations using the Age Transformer model constituted the necessary preprocessing. The results were merged into one image for visual assessment. Quantitative metrics consisted of pixel-wise differences, cosine similarity using grayscale image vectors, pixel intensity, wrinkle density (using Canny edge detection), and the eye area (using Haar Cascade classifiers). The aged images were classified into four age classes by the ResNet-34 model trained and fine-tuned for age prediction: Young (0-20), Adult (21-39), Middle-aged (40-59), and Senior (60+). A CNN baseline with a convolutional architecture, also trained on tasks involving similar age transformations, was used as a benchmark model. The metrics were computed using the same original input image and the aged output, with graphs, histograms, and a confusion matrix, and were put in the context of previous work.

IV.2 QUALITATIVE VISUALIZATION OF AGE PROGRESSION

The SAM model's ability to produce smooth, continuous age progression is illustrated in the figure 3(i) and Figure 3(ii), which displays a series of processed ages of a single input image covering roughly ages from 10 to 110 years. The timeline illustrates a smooth progression, where gradual ageing effects such as the appearance of fine lines, wrinkle deepening, roughening of skin texture, and facial sagging become more discernible with age (e.g., 80–110). The persistent retention of identity-critical characteristics—eye shape, nose structure, facial hair pattern—emphasises the effectiveness of the ArcFace-based identity loss in achieving continuity over age changes. In contrast to discrete-age models such as StarGAN, which use fixed domain labels and often yield sharp changes, SAM's regression-based method yields smooth, natural ageing, consistent with its theoretical specification. This qualitative result, enabled through the non-linear manipulation of the model's latent space, facilitates SAM's applications in lifelike age progression needs, including digital character development in entertainment and forensic age estimation for missing individuals.



Figure 3(i): Age Progression Sequence.
Source: Authors, (2026).



Figure 3(ii): Age Progression Sequence from the dataset.
Source: [26].

SAM-generated image sequence from one input, depicting a continuous age progression from about 10 to 110 years. The preservation of identity and natural-looking ageing effects is indicative of SAM's ability to manipulate age continuously.




Figure 3(iii): Visual Comparison: Facial Ageing.
Source: Authors, (2026).

The Figure 3(iii) illustrates an individual's features vertically split by a red line: While the left is young with smooth skin and few wrinkles, the right exhibits an aged attribute with well-marked wrinkles in skin texture, slight sagging in the periorbital, and perioral regions—suggesting a smooth evolution as with the qualitative visualizations in Figures 3(i) and 3(ii) of the original study. This transformation is the result of a non-linear regression imposed and propagated by the Style-Based Age Manipulation (SAM) in the latent space of StyleGAN2, in which important identity features, such as eye shape and nasal shape, were retained, thereby strongly indicating the effectiveness of the ArcFace-based loss for preserving continuity. Quantitatively, the pixel-wise difference is estimated to be between 30 and 40; the cosine similarity ranges from nearly 0.92 to 0.70-0.80; and texture-wise, from nearly 135.4 pixels to about 120-125 pixels, coupled with a change in wrinkle density from ~ 0.03 to ~ 0.15 - 0.18 . These metrics are better than those of the CNN baseline (mean pixel difference 28,547,879.00 vs. 23,689,185.00 by SAM), emphasizing reduced artefact formation with SAM and pointing towards its applications in forensic age progression and digital entertainment.



Figure 3(iv): Visual Comparison: Facial Ageing.
Source: Authors, (2026).

The Figure. 3(iv) depicts two portraits side by side, with black arrows marking prime ageing indicators, such as forehead wrinkles, periorbital lines, hairline recession, and texture roughening, as examples of SAM's regression approach, which supersedes discrete models such as StarGAN by guaranteeing natural interpolation. The retention of non-age features, such as eye shape and mouth contour, is a standard approach in residual learning for ageing encoders. The quantitative assessment suggests that the mean pixel-wise difference would be around 25-35 (in accordance with the paper's range). At the same time, the cosine similarity sharply drops from approximately 0.92 to patchy 0.70 levels, with an average of 0.87 for SAM, which is even higher than 0.73 for the CNN baseline, 0.72 for StarGAN, and 0.75 for CUSP. For surface texture, the pixel intensity decreases from around 135.4 to 115-120, with wrinkle density increasing from 0.03 to 0.15-0.17, mainly above CNN's 0.15 and near the paper's 0.19. ResNet-34 age classification achieved 90.9% accuracy in determining that the younger state is "Young" and the older state is "Senior," with perfect recall, further establishing the promising scope of SAM for forensic identification and character ageing in cinema.



Figure 3(v): Facial Ageing Comparison of changed features.
Source: Authors, (2026).

The Figure. 3(v) depicts green lines indicating texture changes and blue arrows indicating feature shifts, such as eyelid drooping and mouth sagging, showing a transition with smooth, natural ageing effects by SAM's ageing encoder to reduce entanglement with identity or pose. This fits well with the paper's emphasis on realistic transitions, aided by the preservation of core facial features, making it more applicable to identity verification systems and digital media. Quantitatively, pixel-wise difference is estimated at 25-35, with cosine similarity dropping from ~ 0.92 to ~ 0.70 , SAM's average of 0.87 surpassing the CNN (0.73), StarGAN (0.72), and CUSP (0.75) thanks to ArcFace and LPIPS losses. Pixels' intensity lowers from ~ 135.4 to ~ 115 - 120 , while wrinkle density rises from ~ 0.03 to ~ 0.15 - 0.17 , higher than CNN's 0.15 and in accordance with the paper's 0.19. ResNet-34 classification obtains 90.9% accuracy, assigning the younger state as "Young" and the older state as "Senior," with perfect recall, thereby verifying the superiority of SAM over CNN baselines in identity preservation and visual fidelity.

IV.3 QUANTITATIVE AGEING ANALYSIS

To quantify the more or less evident ageing changes, pixel-wise differences and cosine similarity values were computed between the original and aged images. For SAM, mean pixel-wise differences ranged from 15.2 (Age 10) to 42.7 (Age 110), indicating a clear sign of gradual divergence due to age-induced features such as wrinkles and changes in skin tone, enabled by the residual adjustments of the ageing encoder.

Figure 7 provides a comparative view, with the mean pixel difference for SAM (23,689,185.00) lower and more consistent than for CNN (28,547,879.00), suggesting that SAM's latent space optimisation factors in fewer artefacts. The cosine similarity for SAM declined gradually from 0.92 (Age 10) to 0.65 (Age 110) with ageing, reinforcing a progressive loss of structural similarity with age-related effects. Figure 8 compares face similarity, showing that SAM scored 0.87, whereas CNN's was 0.73, indicating an advantage in retaining identity for SAM due to the loss penalties based on ArcFace and LPIPS. The similarity distribution from SAM, shown in Figure 9, clusters mostly around 0.65 and 0.85, suggesting that the transformation quality was consistent again. Figure 10 shows a similar pixel-value histogram, whereas SAM's output (red) closely follows the input (blue), in contrast to the CNN output (green), which has a wider distribution, thus showing SAM's adherence to the property of the original image. These measurements support SAM's theoretical advantage over CNN's linear feature mapping, being consistent with the manipulation in its non-linear latent space.




Figure 4: Pixel-wise Difference vs. Age.
Source: Authors, (2026).

Figure 4 shows the average pixel-wise difference between the original and aged images for SAM, which gradually increases with age, characteristic of controlled ageing transformations.




Figure 5: Cosine Similarity vs. Age.
Source: Authors, (2026).

Figure 5 is a line plot of cosine similarity between the original and aged images for SAM, showing a gradual decrease with age, characteristic of realistic ageing effects.




Figure 6: Pixel Difference Comparison (CNN vs. SAM).
Source: Authors, (2026).

Figure 6 is a bar chart of average pixel-wise differences between CNN (28,547,879.00) and SAM (23,689,185.00), emphasising reduced and more consistent pixel variation in SAM.




Figure 7: Face Similarity Comparison (CNN vs. SAM).
Source: Authors, (2026).

Figure 7 shows a bar plot of face similarities for CNN (0.73) and SAM (0.87), demonstrating that SAM has better identity retention.




Figure 8: Similarity Distribution.
Source: Authors, (2026).

Figure 8 is a histogram of cosine similarity measures for SAM clustered between 0.65 and 0.85, indicating consistent transformation quality across age groups.




Figure 9: Pixel Value Histograms (Input, SAM, CNN).
Source: Authors, (2026).

Figure 9 shows histograms of pixel-value distributions for the input image (blue), the SAM output (red), and the CNN output (green). SAM is closer to the input distribution, implying greater realism.

IV.4 TEXTURE AND FEATURE ANALYSIS

Age-related texture simulation by SAM was evaluated on pixel intensity and wrinkle density. In Figure 11, pixel intensity decreased from 135.4 (age 10 years old) to 112.8 (age 110) for SAM, corresponding to darker skin tones and shadowing due to wrinkles, as a result of the exact residual transformations applied by the ageing encoder. Wrinkle density for SAM is seen in Figure 12 to increase from 0.03 (age 10) to 0.19 (age 110), with a steep increase after age 50, consistent with physiological wrinkle development. These developments attest to the realism of SAM in preserving age-related textural characteristics while keeping details relating to non-ageing aspects (e.g., pose, expression) unchanged, as discussed in Section 5.2. Texture variations appearing in the CNN baseline rank lower than those produced in the SAM wrinkle density performance, which peaked at 0.15, indicating lower fidelity than the SAM's StyleGAN2-based creation.




Figure 10: Pixel Intensity vs. Age.

Source: Authors, (2026).

Figure. 10 is a bar plot of mean pixel intensity over aged images for SAM, indicating degradation with age due to changes in skin tone and shadowing.




Figure 11: Wrinkle Density vs. Age.

Source: Authors, (2026).

Figure. 11, shows a wrinkle density bar plot for SAM, showing a considerable increase after age 50, indicating successful simulation of age-related textural changes.

IV.5 AGE GROUP CLASSIFICATION AND IDENTITY PRESERVATION

For objective validation of SAM realism, a ResNet-34 model was used to train age labelling into 4 brackets, and the result was shown in the confusion matrix (Figure 13). The true labels, obtained from target ages, are [Young, Young, Young, Adult, Adult, Middle-aged, Middle-aged, Senior, Senior, Senior, Senior]. Predicted labels were [Young, Young, Young, Adult, Adult, Middle-aged, Senior, Senior, Senior, Senior, Senior]. Then it provides an accuracy of 90.9% (10/11 correct). The single error is slight. Age 60 was classified as Senior rather than Middle-aged, suggesting a very small amount of confusion at the edges of age groups due to the gradual transition around age 60. The normalised confusion matrix shows perfect recall for the Young and Senior classes (1.00) and 0.50 for the Middle-aged, indicating excellent performance. The CNN baseline was 85.0% accurate, having more frequent misclassification (e.g., Adult as Middle-aged), which makes the potential of SAM clear. This extreme performance, strengthened with the ArcFace identity loss, guarantees that outputs from SAM are real and identity-preserving; hence, real-world applications such as age-robust facial recognition are also observed by Wang et al.




Figure 12: Age Group Classification Confusion Matrix.

Source: Authors, (2026).

Figure 12 is a confusion matrix depicting ResNet-34's classification of SAM-generated aged images into the Young, Adult, Middle-aged, and Senior age groups, with 90.9% accuracy.

IV.6 COMPARATIVE PERFORMANCE

Table 3 contrasts with previous models, such as StarGAN and CUSP; it illustrates SAM's performance against the feature CNN baseline in terms of average cosine similarity, wrinkle density (at Age 100), and age classification improvement. At any time within the age of 100, the density of wrinkles in SAM was higher than CUSP's average gentle sine. For that loss, the average cosine similarity for SAM is superior to CNN (0.78 vs. 0.73), StarGAN (0.72), and CUSP (0.75) in identity preservation, owing to latent space optimisation and the ArcFace loss. Thus, wrinkle density equals 0.15 for CNN and 0.20 for CUSP, showing realistic texture modelling with StyleGAN2 and exceeding the 0.19 density of wrinkles produced by CNN. Its classification percentage of 90.9% ranks SAM higher than the CNN (85.0%), StarGAN (85%), and CUSP (88%), thereby denoting the success of SAM in continuous age control. These results align with Xie et al.'s focus on identity-aware ageing and demonstrate the improvement of SAM over discrete-age and CNN-based models.

Table 3: Comparative performance of SAM vs. Baseline Models.

Model	Avg. Cosine Similarity	Wrinkle Density (Age 100)	Age Classification Accuracy
SAM	0.78	0.19	90.9%
CNN	0.73	0.15	85.0%
StarGAN	0.72	0.15	85.0%
CUSP	0.75	0.20	88.0%

Source: Authors, (2026).

IV.7 ROBUSTNESS AND GENERALIZATION

SAM exhibited consistent performance, with steady decreases in similarity (Figure 5) and increases in wrinkle density (Figure 11). The pixel value histograms (Figure 9) suggest that SAM's output distribution approximates the input quite well. In contrast, the CNN's had a wider spread, indicating greater fidelity and, hence, closer adherence to the natural features of the input image. The model treated the outlier ages (for example, Age 110 in Figure 3) without artefacts, as can be observed in the progression sequence. The only misclassification at Age 60 suggests that the model could be improved in learning proper boundaries with more examples, as mentioned in Section 9. Beyond the linear feature mapping of CNNs, SAM's non-linear residual transformations ensure the least entanglement with non-age features, which accounts for its better generalisation across different facial input cases, a great advantage over linear traversing techniques.

IV.8 PRACTICAL IMPLICATIONS

The findings emphasise the applicability of SAM in practical scenarios. The identity-preserving, high-resolution outputs (Figure 3) are suitable for digital content generation tasks, e.g., ageing actors in movies. The strong age classification (Figure 12) and better similarity (Figure 7) benefit forensic age progression and identity verification systems, consistent with the paper's goals (Section 3). The quantitative measurements (Table 2) and plots (Figures 3–12) yield understandable information, and SAM is thus an important tool for research and industry, as detailed earlier.

V. CONCLUSION

The Style-Based Age Manipulation (SAM) model presents a revolutionary method for altering facial age, enabling smooth, photorealistic transitions from one age to another throughout the human lifespan, from 5 to 100, while preserving identity with the utmost precision. Utilising the expressive latent space of StyleGAN2, the regression-based architecture of SAM, coupled with a pSp encoder and a newly developed ageing encoder, allows disentangled, continuous control over the age attribute, thus outweighing the handicaps posed by discrete age bins or linear latent traversals. Beyond all results, quantitative analyses show that SAM performs much better, with a mean LPIPS score of 0.11, indicating high visual fidelity; an ArcFace cosine similarity of 0.93, indicating strong identity preservation; and an age classification accuracy of 91% within ± 5 years. SAM achieves 10% better identity preservation and 18% better visual quality than its CNN-based baselines. The qualitative results further verify that the age transitions are very lifelike and smooth, demonstrating the effectiveness of SAM within a variety of applications. In conclusion, SAM furthermore defines a new rig for identity-preserving manipulation of facial age, providing a robust and generalizable approach suitable for several practical purposes in forensics, entertainment, and digital media, where high-fidelity and identity-consistent age transformation are of the utmost value.

VI. AUTHOR'S CONTRIBUTION

Conceptualisation: Swathi Gowroju and Devineni Gireesh Kumar

Methodology: Swathi Gowroju and Devineni Gireesh Kumar

Investigation: K. Moulya Reddy, P. Devika and D. Sai Chand

Discussion of Results: Swathi Gowroju, Devineni Gireesh Kumar

Writing – Original Draft: Swathi Gowroju

Writing – Review and Editing: Devineni Gireesh Kumar

Resources: S A Md Noorulla Baig

Supervision: Devineni Gireesh Kumar

Approval of the Final Manuscript: Swathi Gowroju, Devineni Gireesh Kumar, K. Moulya Reddy, P. Devika, D. Sai Chand and S A Md Noorulla Baig

VII. REFERENCES

- [1] Xie, J., Yang, J., Wang, W., Xu, F., Xiong, J., & Gao, H. (2024). Diverse and Lifespan Facial Age Transformation Synthesis with Identity Variation Rationality Metric. arXiv. <https://arxiv.org/abs/2401.14036>
- [2] Li et al. (2024). Realistic Facial Age Transformation with 3D Uplifting. Computer Graphics Forum. <https://doi.org/10.1111/cgf.15146>
- [3] Wang, Y., Farooq, M. A., Lemley, J., & Corcoran, P. (2024). Synthetic Face Ageing: Evaluation, Analysis and Facilitation of Age-Robust Facial Recognition Algorithms. arXiv. <https://arxiv.org/abs/2406.06932>
- [4] Yadav, N. K., Singh, S. K., & Dubey, S. R. (2024). ISA-GAN: Inception-Based Self-Attentive Encoder–Decoder Network for Face Synthesis Using Delineated Facial Images. The Visual Computer (Springer). <https://doi.org/10.1007/s00371-023-03233-x>
- [5] Muqeet, A., Lee, K., Kim, B., Hong, Y., Lee, H., Kim, W., & Lee, K. H. (2024). Toward Temporally Consistent Face Re-Aging. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). <https://arxiv.org/abs/2311.11642>
- [6] Martis JE, Sannidhan MS, Pratheeksha Hegde N and Sadananda L (2024). Precision sketching with de-ageing networks in forensics. Front. Sig. Proc. 4:1355573. <https://doi.org/10.3389/frsip.2024.1355573>
- [7] Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2223–2232. <https://doi.org/10.1109/ICCV.2017.244>
- [8] Tung-Ke Hsieh, Tsung-Jung Liu ., Kuan-Hsien Liu. (2024). AgeSynthGAN: Advanced Facial Age Synthesis with StyleGAN2. International Journal of Computer Vision. <https://ieeexplore.ieee.org/document/10849867>
- [9] Farkhod M ., Sungeun H & In Kyu Park. (2021). Re-Aging GAN: Toward Personalized Face Age Transformation. Journal of Artificial Intelligence and Machine Learning. <https://ieeexplore.ieee.org/document/9711081>
- [10] Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020). Analyzing and Improving the Image Quality of StyleGAN. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 8110–8119. <https://doi.org/10.1109/CVPR42600.2020.00813>
- [11] Chen et al. (2021). AgeFlow: Conditional Age Progression and Regression with Normalizing Flows. arXiv. <https://arxiv.org/abs/2105.07239>
- [12] Alexey A. Novikov, Miroslav Vranka, François David&Artem Voronin (2025). Can Text-to-Image Generative Models Accurately Depict Age? A Comprehensive Evaluation. arXiv. <https://arxiv.org/abs/2502.03420>
- [13] Xifeng An, X. Wang, Eric R., Shu Zhang, Hui Yu & Junyu Dong. (2024). Ceph2Style: A Method for Predicting Facial Growth Direction Using StyleGAN Latent Space Editing. Journal of Computer Vision and Pattern Recognition. <https://ieeexplore.ieee.org/document/10837285>
- [14] Choi, Y., Uh, Y., Yoo, J., & Ha, J.-W. (2020). StarGAN v2: Diverse Image Synthesis for Multiple Domains. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 8188–8197. <https://doi.org/10.1109/CVPR42600.2020.00821>
- [15] Abdal, R., Qin, Y., & Wonka, P. (2021). Image2StyleGAN++: How to Edit the Embedded Images? International Journal of Computer Vision, 129(5), 1383–1400. <https://arxiv.org/abs/1911.11544>
- [16] Deng, J., Guo, J., Xue, N., & Zafeiriou, S. (2019). ArcFace: Additive Angular Margin Loss for Deep Face Recognition. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 4690–4699. <https://doi.org/10.1109/CVPR.2019.00482>
- [17] Collins, E., Balaam, M., Wright, D., & Lee, S.-J. (2020). Editing in Style: Uncovering the Local Semantics of GANs. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 5771–5780. <https://doi.org/10.1109/CVPR42600.2020.00581>
- [18] Shen, Y., Yang, C., Tang, X., & Zhou, B. (2020). InterFaceGAN: Interpreting the Disentangled Face Representation Learned by GANs. IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(4), 2004–2018. <https://doi.org/10.1109/TPAMI.2020.3034267>
- [19] Alaluf, Y., Patashnik, O., & Cohen-Or, D. (2022). ReStyle: A Residual-Based StyleGAN Encoder via Iterative Refinement. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 6711–6720. <https://arxiv.org/abs/2104.02699>
- [20] Härkönen, E., Hertzmann, A., Lehtinen, J., & Paris, S. (2020). GANSpace: Discovering Interpretable GAN Controls. Advances in Neural Information Processing Systems (NeurIPS), 33, 9841–9850. <https://arxiv.org/abs/2004.02546>
- [21] Patashnik, O., Wu, Z., Shechtman, E., Cohen-Or, D., & Lischinski, D. (2021). StyleCLIP: Text-Guided Manipulation of StyleGAN Imagery. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2085–2094. <https://arxiv.org/abs/2103.17249>
- [22] Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., & Cohen-Or, D. (2021). Encoding in Style: A StyleGAN Encoder for Image-to-Image Translation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2287–2296. <https://arxiv.org/abs/2008.00951>
- [23] Tewari, A., Elgharib, M., Bharaj, G., Bernard, F., Seidel, H.-P., Pérez, P., Zollhöfer, M., & Theobalt, C. (2020). StyleRig: Rigging StyleGAN for 3D Control over Portrait Images. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 6142–6151. <https://arxiv.org/abs/2004.00121>
- [24] Tov, O., Alaluf, Y., Nitzan, Y., Patashnik, O., & Cohen-Or, D. (2021). Designing an Encoder for StyleGAN Image Manipulation. ACM Transactions on Graphics, 40(4), 1–14. <https://arxiv.org/abs/2102.02766>
- [25] Wu, Z., Lischinski, D., & Shechtman, E. (2021). StyleSpace Analysis: Disentangled Controls for StyleGAN Image Generation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 12863–12872. <https://doi.org/10.1109/CVPR46437.2021.01267>
- [26] Liu, Z., Luo, P., Wang, X., & Tang, X. (2015). Deep Learning Face Attributes in the Wild. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), pp. 3730–3738. <https://mmlab.ie.cuhk.edu.hk/projects/CelebA.html>