

PEGASUS: Personalized Generative 3D Avatars with Composable Attributes

Hyunsoo Cha Byungjun Kim Hanbyul Joo
 Seoul National University

243stephen@snu.ac.kr byungjun.kim@snu.ac.kr hbjoo@snu.ac.kr

<https://snuvclab.github.io/pegasus/>

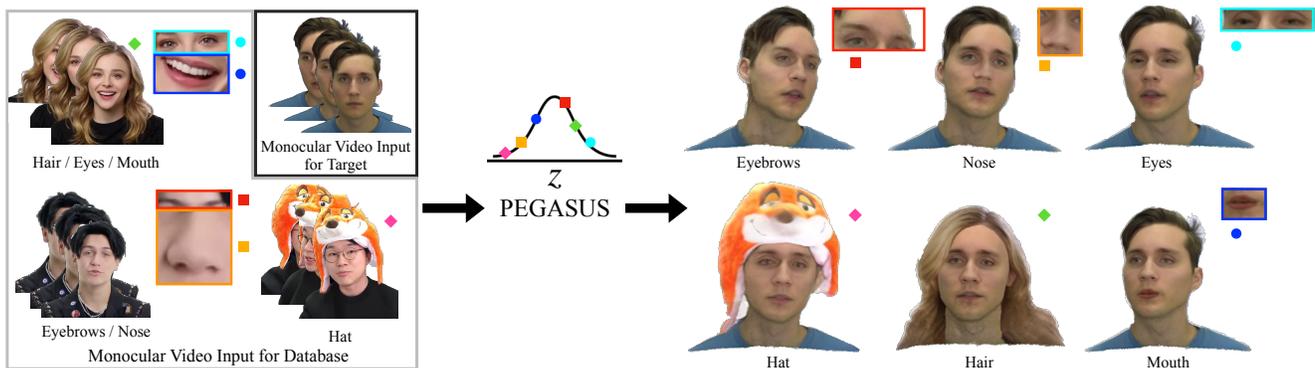


Figure 1. PEGASUS. Our method builds a personalized generative 3D face avatar from monocular video sources.

Abstract

We present PEGASUS, a method for constructing a personalized generative 3D face avatar from monocular video sources. Our generative 3D avatar enables disentangled controls to selectively alter the facial attributes (e.g., hair or nose) while preserving the identity. Our approach consists of two stages: synthetic database generation and constructing a personalized generative avatar. We generate a synthetic video collection of the target identity with varying facial attributes, where the videos are synthesized by borrowing the attributes from monocular videos of diverse identities. Then, we build a person-specific generative 3D avatar that can modify its attributes continuously while preserving its identity. Through extensive experiments, we demonstrate that our method of generating a synthetic database and creating a 3D generative avatar is the most effective in preserving identity while achieving high realism. Subsequently, we introduce a zero-shot approach to achieve the same goal of generative modeling more efficiently by leveraging a previously constructed personalized generative model.

1. Introduction

Building a personalized 3D avatar for representing an individual in virtual spaces can bring significant advancements

in the field of AR/VR and applications within the metaverse. Importantly, the method should be user-friendly to allow novices to build their avatars without the need for complex capture systems. It should also offer a high level of realism, depicting the fine-grained details of the individual’s geometry and appearance, and, importantly, the avatar should be animatable to mirror the user’s facial expressions in the virtual space. However, the 3D avatar does not need to maintain an exact replica of the user’s single appearance, as users may prefer to alter their avatars. This includes modifications of changing hairstyles, adding accessories like hats, or even altering facial parts to give the avatar a more aesthetically pleasing look, such as adopting the appearance of celebrities.

Recent technologies make it possible to build high-quality 3D face avatars for general users from monocular video inputs only [2, 14, 16, 17, 46, 47]. By leveraging parametric morphable face models [3, 29], these approaches produce realistic animatable human avatars from sparse monocular videos that capture naturally moving facial images, by fusing observed cues into a canonical space. However, the previous approaches mainly focus on creating the exact replica from the input videos, without providing the functionality to alter the subparts of the avatars, such as hairstyles or nose. As an alternative direction, generative models in producing realistic faces have been studied in the 2D field, by producing 2D

human faces with diverse appearance changes and facial expression changes [37, 39, 49]. 3D-aware generative models leveraging the pre-trained 2D generative models are also presented for generative face modeling in 3D [6, 7]. While this approach shows realistic faces, they are not fully animatable, lacking explicit mapping to the 3D morphable models, and, thus, it is difficult to reenact the facial expressions from the target or allow viewpoint changes while keeping the identity.

In this work, we present PEGASUS, a method to build a *personalized generative 3D avatar* from monocular video inputs. In contrast to the previous work [46, 47], our 3D avatar enables compositional controls of facial attributes, where users can make alterations for desired facial attributes such as hair, nose, or accessories, as shown in Fig. 1, while preserving the identity of the target person. The control can be performed by changing the disentangled latent codes defined in a continuous latent space. Our personalized generative 3D avatar is constructed from the monocular video of the target individual. Importantly, to learn the possible variations of each facial attribute, we leverage other available monocular videos from arbitrary individuals, where our personalized generative models can automatically learn continuous disentangled latent spaces of facial attributes.

However, there exist significant challenges in consolidating the monocular videos from multiple individuals into a personalized generative model construction for the target individual. Building a model with videos from many individuals often fails to preserve the fine-grained appearance details of the target individual, and, more critically, changing the latent space can lead to changes in the entire facial appearance, rather than selectively altering the desired subpart. As a solution, we present an approach by synthesizing part-swapped videos of the target individual by replacing a specific facial part with the one from other individuals, as shown in Fig. 4. Built with diverse part-swapped videos, our generative 3D avatar, PEGASUS, can preserve high-quality details for the target individuals, while equipped with the generative power to selectively alter each facial part. While our generative 3D avatar already shows satisfactory performance, it involves the time-consuming process of constructing a set of part-swapped videos. As a more rapid and efficient solution, we further introduce an approach that achieves the same objectives through *zero-shot part transfer*, leveraging previously constructed personalized generative models. Through several experiments, we demonstrate the superior performance of our approach when compared to alternative methods.

Our contributions are summarized as follows: (1) the first method to build personalized generative 3D avatars from monocular video sources; (2) disentangled controllability to selectively alter a subpart or multiple parts of the 3D faces of the target individual; and (3) the 3D part transfer approach to efficiently implement personalized generative models without additional training.

2. Related Work

3D Facial Avatar Reconstruction. To deal with the inherent diversity and dynamics of human faces, 3D parametric face models have been proposed to represent 3D facial changes via a set of parameters that model variations of the shapes, poses, and expressions of faces [3, 5, 29]. DECA [11] proposes a method for regressing FLAME [29] parameters from monocular images, which enables 3D facial avatar reconstruction without a 3D scan setup. With the advancements of neural rendering [32], several neural rendering-based 3D facial avatar reconstruction approaches were proposed to overcome the limited facial details of the parametric model [1, 9, 11, 14, 16, 35, 40, 46, 47, 52]. IMAvatar [46] introduces the dynamic 3D morphable face model as an implicit representation with canonicalization of the head deformation and facial expression based on SNARF [8]. PointAvatar [47] proposes a deformable point-based representation to reconstruct high-frequency details from monocular video with efficient rendering.

Face Editing in 2D/3D. GAN [15] has made significant contributions to producing natural and high-quality face editing [6, 7, 25, 37, 39, 42, 44, 48–50]. For instance, Barbershop [49] edits the hairstyle of the target face to match the source appearance by manipulating the latent space that represents the feature’s spatial location and appearance. With the rise of the diffusion model [21, 38], methods for editing faces or scenes based on diffusion have been proposed [4, 19, 31]. For instance, Instruct-Pix2Pix [4] introduces a method for editing an input image with text instructions by finetuning a pretrained diffusion model [34] using a generated image editing dataset [19] through supervised learning. Following the emergence of 2D diffusion models, several approaches leverage the pretrained diffusion models to edit 3D scenes from the text prompts [18, 36, 51]. For instance, Instruct-NeRF2NeRF [18] proposes an iterative optimization process for editing a pretrained NeRF scene [32]. However, these approaches struggle to preserve the identity of the scene and require additional optimization.

Compositional Modeling for 3D Face Avatar. Several methods propose to edit implicit representations leveraging learnable or pretrained latent spaces [20, 22, 24, 41]. However, due to the difficulty in editing, recent approaches introduce decoupled representations for garments or attributes [12, 13, 20, 26, 28]. MEGANE [28] attaches reconstructed eyeglasses from 3D scans to volumetric primitive 3D avatars, optimizing deformation for avatar-specific adjustments. SCARF and DELTA [12, 13] create 3D avatars using a hybrid representation, enabling the transfer of garments or hair without additional optimization. However, these approaches are limited to specific categories or attributes, such as eyeglasses and bags, and the synthesized result avatars are not natural. Furthermore, they lack the ability to change or generate facial attributes continuously.

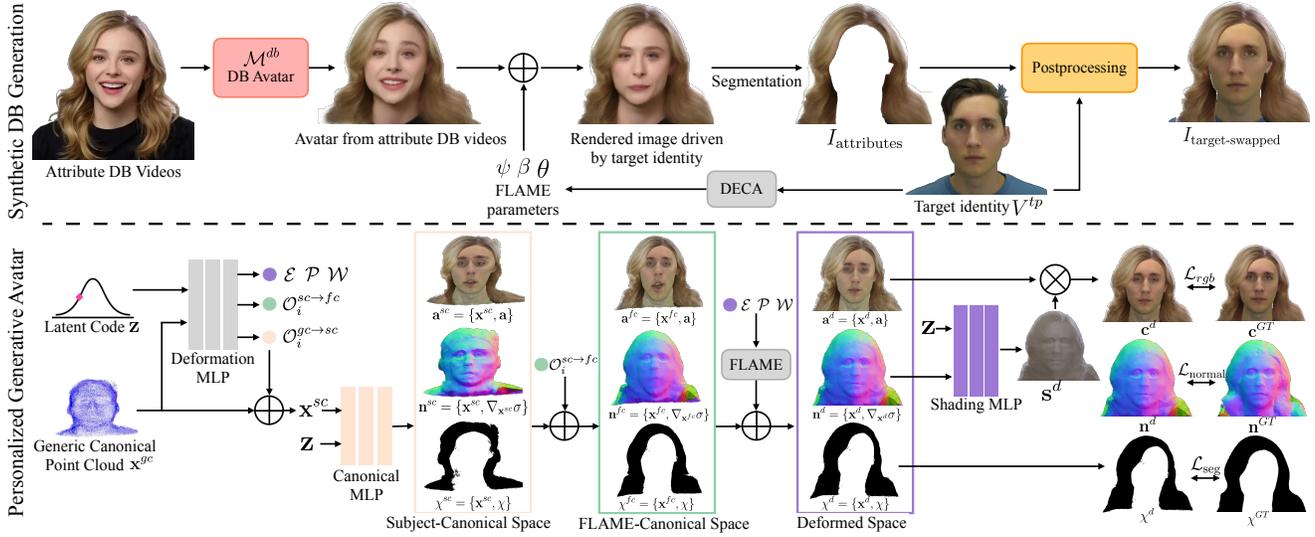


Figure 2. **Method Overview.** Our approach consists of two main components: synthetic database (DB) generation and a personalized generative 3D avatar model. Initially, we build a synthetic DB via face part swapping from the attribute DB videos. For the generation of the synthetic DB, we propose the method through post-processing and attribute alignment leveraging FLAME parameters. Subsequently, we train our model utilizing the synthetic DB that contains the same target identity with varying attributes. Our model infers the 3D point locations in the deformed space \mathbf{x}^d , normal \mathbf{n}^d , shading \mathbf{s}^d , point segment cues χ^d , and the albedo \mathbf{a}^d for each queried canonical point \mathbf{x}^{gc} , conditioned by the latent code \mathbf{z} .

3. Preliminaries: PointAvatar [47]

Our approach extends PointAvatar, which reconstructs the 3D avatar of a single identity from a monocular video, to a personalized generative avatar with controllable facial attributes. PointAvatar represents the target avatar via the initial canonical learnable point representations $P_c = \{x_i^c\}_{i=\{1 \dots N\}}$, where $x_i^c \in \mathbb{R}^3$ represents i -th learnable point defined in the canonical space (denoted as the superscript c). By estimating the offset value $\mathcal{O}_i^{c \rightarrow fc}$ from a trained MLP, the canonical points are deformed into the *FLAME-canonical space* (denoted as fc) as: $\mathbf{x}_i^{fc} = \mathbf{x}_i^c + \mathcal{O}_i^{c \rightarrow fc}$. Subsequently, the points are deformed into the posed space as leveraging FLAME model [29]:

$$\mathbf{x}^{d-} = \mathbf{x}^{fc} + B_P(\boldsymbol{\theta}; \mathcal{P}) + B_E(\boldsymbol{\psi}; \mathcal{E}) \quad (1)$$

$$\mathbf{x}^d = \text{LBS}(\mathbf{x}^{d-}, \mathbf{J}(\boldsymbol{\psi}), \theta, \mathcal{W}), \quad (2)$$

where \mathbf{x}^{d-} denotes the point after applying the blendshapes and before applying transformation via linear blend skinning (LBS). $\boldsymbol{\psi}$, θ , β are the expression, pose, and shape parameters of the FLAME model, respectively, for animating the avatar, and \mathcal{E} , \mathcal{P} , and \mathcal{W} are the expression blendshapes, pose blendshapes, and LBS weights, respectively, which are estimated by an MLP. The normal of each point \mathbf{n}_c is defined as a signed distance field (SDF), which is the canonical network's output as follows: $\mathbf{n}_c = \nabla_{\mathbf{x}_c} \text{SDF}(\mathbf{x}_c)$. The normal of the deformation space \mathbf{n}_d is represented by a deformation network which deforms the canonical point set P_c to the deformed point set $P_d = \{\mathbf{x}_d^i\}$. The point deformation is

fully differentiable, so it can define the normal deformation as follows:

$$\mathbf{n}_d = l \mathbf{n}_c \left(\frac{\partial \mathbf{x}_d}{\partial \mathbf{x}_c} \right)^{-1}, \quad (3)$$

where l denotes the normalizing factor, which ensures the output of normal value should be the unit length. The RGB of a point is represented by $c_d = s_d \circ a$, the Hadamard Product of the shading s_d , and albedo a .

4. Method

4.1. Personalized Generative Avatar Model

Our generative avatar model takes a latent code $\mathbf{z} \in \mathbb{R}^{(D+1) \times d}$ and FLAME parameters β , θ , and $\boldsymbol{\psi}$ as inputs. The latent code is the concatenation of the D part-wise latent codes $\{\mathbf{z}^j\}_{j=0 \dots D}$, where each part-wise latent code $\mathbf{z}^j \in \mathbb{R}^d$ controls the identity of the humans or the subpart such as hair and nose. We treat that \mathbf{z}^0 controls the overall identity variations while changing other codes $\mathbf{z}^{j \neq 0}$ varies only the subparts of the face, preserving the same identity represented by \mathbf{z}^0 . By changing FLAME parameters, we can animate the avatars to have varying face poses and expressions. The shape parameter of FLAME β also affects the overall coarse shape of the avatar, and we assume the parameter is fixed for the same individual with the same \mathbf{z}^0 . By extending the PointAvatar [47], our avatar model is represented by a set of *generic* (or person-agnostic) canonical point $P^{gc} = \{\mathbf{x}_i^{gc}\}_{i=\{1, \dots, N\}}$. To this end, our model, \mathcal{M}_ϕ , infers the 3D point locations in the deformed space \mathbf{x}_i^d ,

normal vector $\mathbf{n}_i \in \mathbb{R}^3$, and the albedo color $\mathbf{a}_i \in \mathbb{R}^3$ for each queried canonical point as:

$$\mathcal{M}_\phi(\mathbf{x}_i^{gc}, \mathbf{z}, \beta, \boldsymbol{\theta}, \boldsymbol{\psi}) = \{\mathbf{x}_i^d, \mathbf{n}^d, \mathbf{a}_i\}, \quad (4)$$

where \mathbf{x}_i^d represents the 3D point after applying identity and appearance variations controlled by \mathbf{z} , as well as the facial pose and expression changes by FLAME parameters. Fig. 2 represents an overview of PEGASUS.

In contrast to the original PointAvatar, which represents a single identity only, we train a single avatar model to represent multiple face appearances, where appearance can vary by changing disentangled latent codes \mathbf{z} . We tackle this challenging problem by introducing the generic canonical space, which is person-agnostic. While our model ideally expresses a range of identities, we observe that representing extremely diverse individuals with a single implicit model often results in blurry avatars, as demonstrated in our ablation studies. Yet, we demonstrate that our model can successfully achieve the goal of a personalized generative avatar model, allowing face part variation while preserving the same identity. Importantly, in order to build the personalized generative avatar model, we present a way to synthesize the dataset of the target individual via part-swapping, described in Sec. 4.2.

Multi-stage Canonical Spaces and Point Deformation.

While the original PointAvatar considers two-staged deformation (canonical, FLAME-canonical, and deformed space), we consider one more stage, resulting in generic canonical (gc), subject-specific canonical (sc), subject-specific FLAME-canonical (fc), and deformed space (d). We empirically find that introducing the generic canonical space enables us to avoid bad local minima in training the model with multiple face appearances, which is similar to PointAvatar, and to enhance identity preservation while varying the attributes by randomly sampled latent code. See more details in Supp. Mat.

The generic canonical space and the point locations defined in this space $P^{gc} = \{\mathbf{x}_i^{gc}\}_{i=1\dots N}$ are shared among all identities. We first map the points \mathbf{x}_i^{gc} from the generic canonical space into the subject-specific canonical space $P^{sc} = \{\mathbf{x}_i^{sc}\}_{i=1\dots N}$ by adding point offsets $\mathcal{O}_i^{gc \rightarrow sc}$ that are conditioned by latent code \mathbf{z} . Subsequently, we map the points in the subject-specific canonical space into the FLAME-canonical space via another point offset $\mathcal{O}_i^{sc \rightarrow fc}$, similar to the PointAvatars. That is,

$$\mathbf{x}_i^{sc} = \mathbf{x}_i^{gc} + \mathcal{O}_i^{gc \rightarrow sc} \quad (5)$$

$$\mathbf{x}_i^{fc} = \mathbf{x}_i^{sc} + \mathcal{O}_i^{sc \rightarrow fc}, \quad (6)$$

where $\mathcal{O}_i^{gc \rightarrow sc}$ and $\mathcal{O}_i^{sc \rightarrow fc}$ are inferred from the learned deformation MLP model. Intuitively, our subject-specific canonical space is equivalent to the ‘‘canonical space’’ of PointAvatars, where we introduced one more prior stage to handle multiple identities.



Figure 3. **DB Avatar.** We create deformable avatar models from the attribute DB videos, which are monocular RGB inputs. We show some examples of our collection of avatars from attribute DB videos with the same FLAME parameters.

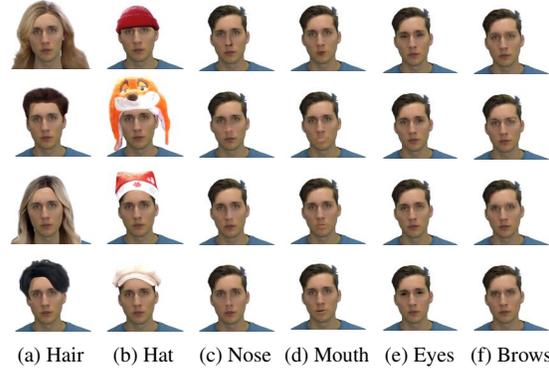


Figure 4. **Part-Swapped Videos of the Target Individual.** Some examples of the synthetic DB created through part-swapping. Our synthetic DB includes a variety of hair, hats, eyes, noses, mouths, and eyebrows.

As in PointAvatars, we use a coordinate-based MLP to infer deformation offsets, blendshapes, and LBS weights:

$$MLP(\mathbf{z}, \mathbf{x}_i^{gc}) = \{\mathcal{O}_i^{gc \rightarrow sc}, \mathcal{O}_i^{sc \rightarrow fc}, \mathcal{E}, \mathcal{P}, \mathcal{W}\}. \quad (7)$$

The deformed point is then computed as:

$$\mathbf{x}^{d-} = \mathbf{x}^{fc} + B_S(\beta; \mathcal{S}) + B_P(\boldsymbol{\theta}; \mathcal{P}) + B_E(\boldsymbol{\psi}; \mathcal{E}) \quad (8)$$

$$\mathbf{x}^d = \text{LBS}(\mathbf{x}^{d-}, \mathbf{J}(\boldsymbol{\psi}), \theta, \mathcal{W}). \quad (9)$$

Different from PointAvatar, we leverage the shape blendshapes basis B_S of the FLAME, allowing us to change the coarse shape of the avatar by controlling the shape parameter β of the FLAME, which is useful for building our synthetic DB to enable better face alignments described in Sec. 4.2.

Canonical Representations. We utilize an MLP to infer the SDF value $\sigma_i \in \mathbb{R}$, albedo $\mathbf{a}_i \in \mathbb{R}^3$, shading $\mathbf{s}_i \in \mathbb{R}^3$, point segment cues $\chi_i \in [0, 1]$ for the i -th point at the subject canonical space \mathbf{x}_i^{sc} :

$$MLP(\mathbf{z}, \mathbf{x}_i^{sc}) = \{\sigma_i, \mathbf{a}_i, \chi_i\}. \quad (10)$$

Note that we consider these cues on the subject canonical space \mathbf{x}_i^{sc} , rather than the generic canonical space since we empirically find inferring it in the generic canonical space

suffers from local minima issue. Similar to PointAvatar, the SDF cues are utilized to infer surface normal in the subject canonical space \mathbf{n}_{sc} and the ones in the deformed space can be computed as in Eq. (3). Note that, different from PointAvatar, the cues in the canonical space are also conditioned by latent code \mathbf{z} , allowing for varying by controlling latent codes for part appearance changes. Furthermore, we additionally include the binary segmentation cues χ_i to estimate a “synthesis” part in the current identity represented by \mathbf{z} , which is used in our *Zero Shot Transfer* approach in Sec. 5.

Comparison over PointAvatar. The major difference from the PointAvatar is the use of latent codes \mathbf{z} to enable the single model can handle varying appearance changes. For this purpose, we modify the model, including a generic canonical stage and injecting the \mathbf{z} into the submodules. We also made several modifications, including (1) the beta controlling part in Eq. (8), which is essential for fitting the subject and face attribute when generating a synthetic database, (2) inferring the segmentation mask for the usage in *Zero Shot Transfer*.

4.2. Synthetic DB Generation via Part Swapping

We aim to build our personalized generative model to preserve the target human identity, while allowing changing facial attributes, such as hair, nose, or wearing a hat. To learn such a model, we need the videos of the target human with all such variations, which is not available in practice. We present a solution to synthesize such variations from other video sources by swapping a face subpart of the target identity person with others. Examples are shown in Fig. 4. We collect a set of monocular videos, denoted as $V^{db} = \{V_i, p_i\}_{i=1 \dots K}$ from various individuals to model various types of facial attribute variations. For each video V_i , we determine the target facial attribute $p_i \in \mathbf{P}$ which we want to use for the swapping, where $\mathbf{P} = \{\text{hair, nose, hat, eyes, eyebrows, mouth}\}$.

For each monocular video V_i from the facial attribute DB, we build a personalized avatar \mathcal{M}_i^{db} using our modified avatar generation module which removes the subject-canonical space with only the single video identity as shown in Fig. 3. Here we set the identity latent code \mathbf{z}_0 as learnable while setting other parts accordingly. Once built, the avatar is animatable following the FLAME parameters.

Face Part Swapping. We denote the input video of the target person as V^{tp} . The goal of our face part swapping is to replace the facial attributes V^{tp} with the one using the person appeared in by i -th attribute video V_i . Since both videos have different poses, viewpoints, and facial expressions of different individuals, such replacement is non-trivial in 2D video space. Our idea is to leverage the animatable avatar model \mathcal{M}_i^{db} constructed from V_i to render the facial attribute aligned into the target identity’s videos, V^{tp} . This can be performed by inputting the FLAME parameters and camera parameters obtained from V^{tp} into \mathcal{M}_i^{db} and by rendering only the necessary attribute region with blending. To choose



Figure 5. **Zero Shot Transfer.** PEGASUS generates high-quality and natural appearances through zero-shot transfer.

the selected attribute regions specified by the corresponding attribute $p_i \in \mathbf{P}$, we use an off-the-shelf face part segmentation model [43] to obtain the mask of the desired target attribute p_i .

Then, we can synthesize the attribute part into the target human videos V^{tp} as follows:

$$I_{i\text{-th attributes}} = \mathcal{R}(\mathcal{M}_i^{db}(\theta_{tp}, \beta_{tp}, \psi_{tp}))$$

$$I_{\text{target-swapped}} = \mathbf{1}_{tp} \cdot I_{tp} + \mathbf{1}_{i\text{-th attributes}} \cdot \mathcal{B}(I_{i\text{-th attributes}}),$$

where \mathcal{R} denotes the rendering function from the avatar \mathcal{M}_i^{db} with the FLAME parameters obtained from V^{tp} . $\mathbf{1}_{tp}$ is the segmentation mask to select the target subject regions excluding the attribute parts, and $\mathbf{1}_{i\text{-th attributes}}$ is segmentation for the attributed part of from \mathcal{M}_i^{db} ’s rendering, respectively. Note that we use the shape parameter of the target human β_{subject} on \mathcal{M}_i^{db} as an input to make better alignment into the target identity, which was the motivation for introducing the shape parameter β in building our avatar, different from the original PointAvatar model. \mathcal{B} denotes the blending function, where we use Poisson Blending [33] to reduce artifacts. We further perform post-processing to enhance the quality of part-swapped images using OpenCV’s dilate and erode function to remove holes. As a special preprocessing for hair-swapping, it is empirically advantageous to synthesize the target person’s hair into a bald head before the blending, where we leverage Stable Diffusion [34] with auto-generated mask images. See Supp. Mat. for details.

We denote \hat{V}_i^{tp} as the part-swapped videos by i -th attribute DB identity. Examples are shown in Fig. 4. Note that the resulting videos contain the same target identity with varying attributes via synthesis, which we use to build our personalized generative models.

4.3. Learning for Personalized Generative Model

Latent Code Setting. We train our model by using V^{tp} and synthesized videos $\{\hat{V}_i^{tp}\}_{i=1 \dots K}$. For each video, we set the latent code $\mathbf{z} = \{\mathbf{z}^p\}_{p=0, \dots, D}$ according to the attribute types. Specifically, we use the same shared learnable identity latent code \mathbf{z}_0 for all videos, given that the videos are for the same identity. If a video is about the variation of p -th attribute category, where $p \in \mathbf{P}$, we assign a separate learnable latent

code for that part \mathbf{z}^p , by keeping other latent code parts shared. With this setup, we allow the model can have the latent codes in a disentangled manner so that each attribute code part can represent the corresponding facial subparts.

Loss Function. We follow the PointAvatar [47] to define loss functions. The total loss is as follows:

$$\mathcal{L} = \lambda_{\text{rgb}}\mathcal{L}_{\text{rgb}} + \lambda_{\text{mask}}\mathcal{L}_{\text{mask}} + \lambda_{\text{FLAME}}\mathcal{L}_{\text{FLAME}} + \lambda_{\text{vgg}}\mathcal{L}_{\text{vgg}} + \lambda_{\text{normal}}\mathcal{L}_{\text{normal}} + \lambda_{\text{seg}}\mathcal{L}_{\text{seg}} + \lambda_{\text{z reg}}\mathcal{L}_{\text{z reg}},$$

where \mathcal{L}_{rgb} , $\mathcal{L}_{\text{mask}}$, $\mathcal{L}_{\text{FLAME}}$ penalize RGB, mask, FLAME parameter differences respectively. \mathcal{L}_{vgg} is based on the VGG feature to enhance the rendered image quality. Different from previous work [46, 47], we also include three more losses, $\mathcal{L}_{\text{normal}}$, \mathcal{L}_{seg} and $\mathcal{L}_{\text{z reg}}$. We adopt the normal loss as follows: $\mathcal{L}_{\text{normal}} = \|\mathbf{n} - \mathbf{n}^d\|$, and we empirically find its advantage in producing better-quality avatars. We generate the pseudo ground truth normal \mathbf{n} from the V^{tp} and the avatar trained with a single identity of each V^{db} . We also include segmentation loss to predict facial attribute categories χ_i per each point. See more details in Supp. Mat.

Training Strategy. We train PEGASUS in a coarse-to-fine manner. First, following PointAvatar, we upsample the number of points and reduce the radii of the points during the training with the constant period of epochs. Second, we train our model in the two-step strategy. Initially, we train our model using the target individual V^{tp} , which is no part swapped on the face, with the latent codes \mathbf{z} until the beginning of the training. Subsequently, we use all of the part-swapped videos \hat{V}_i^{tp} until the end of training. Check the details in Supp. Mat.

5. Generative Avatar via Zero-Shot Transfer

We present an alternative method to efficiently achieve the goal of a personalized generative avatar without producing part-swapped synthesized videos. Our core idea is based on the assumption that we already have the previously constructed personalized avatar model for an identity (denoted as the *source human*), \mathcal{M}_ϕ , with the functionally to control the face attribute variations. Given the new identity’s video (denoted as the *target human*), we first train our generative avatar architecture with the single video of the target human, resulting in \mathcal{M}_{th} . Then, we aim to achieve the same goal of the personalized avatar for the target human, by fusing the controlled attributed part of \mathcal{M}_ϕ and the remaining part \mathcal{M}_{th} , which we call a “zero-shot model”. Especially, given the FLAME parameters and input latent codes inputs, we can drive both models as:

$$\mathcal{M}_\phi(\mathbf{x}^{gc}, \mathbf{z}, \beta, \theta, \psi) = \{\mathbf{x}_\phi^d, \mathbf{n}_\phi^d, \mathbf{a}_\phi, \chi_\phi\}, \quad (11)$$

$$\mathcal{M}_{th}(\mathbf{x}^{gc}, \mathbf{z}, \beta, \theta, \psi) = \{\mathbf{x}_{th}^d, \mathbf{n}_{th}^d, \mathbf{a}_{th}, \chi_{th}\}. \quad (12)$$

The final version of the avatar is constructed by combining the subsets of point clouds from both avatars, using the estimated segmentation masks, χ_ϕ and χ_{th} :



Figure 6. **Single Part-Swap Avatar on Hair.** Our synthesis method creates a photo-realistic avatar with a hairstyle that is accurately transferred.

$$P_{\text{naive}} = \{\mathbf{x}_\phi^{d,i}\}_{\chi_\phi^i=1} \cup \{\mathbf{x}_{th}^{d,i}\}_{\chi_{th}^i=0}, \quad (13)$$

where χ_ϕ^i and χ_{th}^i are the segmentation masks of the face attribute we currently try to control via \mathbf{z} . Intuitively, we replace the points of the desired facial attribute of the novel target human with those of the pretrained generative avatar. While we find this naive composition is already compelling, we observe that there exists a gap between the fused parts. To enhance the quality, we further perform an additional optimization processing to better alignment, with color blending. See Supp. Mat. for the post-processing. Examples of our zero-shot modeling are shown in Fig. 5.

6. Experiments

Datasets. As the attribute database, we collect publicly available 109 videos from the Internet, and build their individual DB avatar model \mathcal{M}^{db} as shown in Fig. 3. For the target person V^{tp} used of the personalized generative avatar, we select the publicly available videos from NerFACE [14], and the individuals are shown in Fig. 4 and Fig. 6. To reenact the reconstructed avatar into unseen facial poses and expressions, we extract FLAME parameters using DECA [11] from our own monocular video with diverse facial orientations and expressions.

6.1. Part-Swapping Comparison with Baselines

Given that we are the first to build a personalized generative model, there is no direct competitor to compare the full generative functionality. Thus, we consider a sub-problem of building an animatable 3D avatar by transferring a facial attitude from another video source. While the resulting output is not a generative model due to its limitation of producing unseen attributes, one can use this strategy to alter parts of the

face, assuming a large number of attribute source videos are available. In this evaluation, we only consider hairstyles as our attribute and consider 5 videos with different hairstyles. Examples are shown in Fig. 6.

Baselines. We consider possible alternative approaches to building the 3D avatar of the target individual with the hair from another video.

DELTA [13]: DELTA achieves the transfer of hairstyles from a source to a target by employing a hybrid approach that combines both explicit and implicit representations. The major goal of DELTA is aligned with this sub-problem test, while it does not have generative functionality.

E4S [30] + PointAvatar (E4S+PA): E4S employs GAN inversion for the face swapping. As a way of building a 3D avatar, we first replace the hair of the target individual in 2D spaces on all image frames via E4S. Then, we apply the original version of PointAvatar to make it into a 3D avatar model. Note that the GAN-based method does not guarantee the view consistency on the synthesized images, resulting in blurry 3D avatar construction.

Custom Diffusion [27] + PointAvatar (CD+PA): Similar to the E4S+PA, we can apply the Custom Diffusion model as a tool to produce hair-changed 2D images for the target individual, conditioned by the hair-style of other video source. Then, we apply the original PointAvatar.

Ours_{swap} + PointAvatar (Ours_{swap} + PA): We also include a simplified version of ours as a baseline, where we produce the part-swapped 2D videos (described in Sec. 4.2) for each hairstyle transfer and apply PointAvatar.

Ours_{person-gen} and Ours_{zero-shot}: We show the performance of our generative models using the latent codes corresponding to the target hairstyles. Note our models can produce not just various hairs, but all other attribute styles.

Metrics. After we build 3D avatars of the target individual by transferring the hairstyle from video sources, we apply unseen facial expressions and head orientation to visualize the avatar in diverse novel poses and render them into images. For the comparison, we consider both the naturalness or the 3D avatar and identity preservation of the target individual. We use two metrics, Fréchet Inception Distance (FID) and Kernel Inception Distance (KID), to evaluate the naturalness of the rendering of the produced 3D avatar. In computing FID and KID, we compare the distributions of rendered outputs of the 3D avatars with the background matted FFHQ [23]. To quantify whether the output 3D avatars preserve the original identity of the target human, we include ArcFace [10] metric. Here, we compare the rendering of the edited version with the rendering of the non-edit PointAvatar with the same unseen face pose.

Results. We show the quantitative comparison in Tab. 1 and example results in Fig. 6. As shown in the table, the 3D avatar produced by our face-swap *Ours_{swap}+PA* achieves the best metrics at FID and ArcFace, showing better naturalness

Method	Naturalness		Identity
	FID↓	KID↓	ArcFace↑
CD + PA	181.60	0.1367	0.6691
E4S + PA	176.64	0.1416	0.5701
DELTA	198.40	0.1797	0.6732
Ours _{swap} +PA	169.54	0.1406	0.7179
Ours _{person-gen}	190.10	0.1696	0.6883
Ours _{zero-shot}	191.47	0.1881	0.7792

Table 1. **Quantitative Comparison.** The synthesis method (upper rows) and full model of hair category (lower rows).

Method	Naturalness		Identity
	FID↓	KID↓	ArcFace↑
Ours _{no synthesis, latent swap}	231.62	0.2630	0.6285
Ours _{no synthesis, latent interp.}	240.17	0.2482	0.4722
Ours _{synthesis, latent interp.}	206.87	0.1839	0.8127

Table 2. **Evaluating Generative Performance.** Quantitative comparison by producing appearance via latent code interpolation.

while keeping the identity of the target individual. Although the custom diffusion-based output *CD+PA* shows the better result in the KID metric, it changes the identity significantly, resulting in low performance in the ArcFace metric. Our full generative model *Ours_{person-gen}* also shows convincing performance even though the model is much more generic and trained to express diverse variations. It outperforms all other baseline methods in preserving identity while showing comparable naturalness. Our zero-shot generative model *Ours_{zero-shot}* shows the best identity-preserving performance because its face part is identical to the non-edited PA while transferring the hair part from *Ours_{person-gen}*.

6.2. Evaluating Generative Performance

We also compare the generative performance of our models. As the baseline, we consider the scenario of using entire videos, including target individual V^{tp} and face attribute videos V^{db} into the generative model without our facial part-swap approach. Once trained, we check the unseen appearances by interpolating the latent codes of two seen samples during training. However, we consider two ways of interpolation: (1) naive interpolation between the latent codes ($\mathbf{z}_A, \mathbf{z}_B$) of two original videos (latent interpolation), (2) via latent code swapping by keeping the target individual’s latent code $\{z_p\}_{p \neq i}$ and other sources’ attribute latent code $\{z_q\}_{q=i}$ (latent swapping), $\mathbf{z} = \{z_p, z_q\}_{p \neq i, q=i, i \in [0, n(\mathbf{P})]}$. For quantitative evaluation, we use the same metric as Sec. 6.1 to measure the naturalness and identity preserving. The quantitative result is shown in Tab. 2, and example qualitative results are shown in Fig. 7. The outputs of our model show compelling performance in producing realistic

\mathcal{O}_{gc}	\mathcal{O}_{sc}	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
×	×	20.92	0.9059	0.1351
×	✓	21.55	0.9033	0.1292
✓	✓	21.75	0.9059	0.1291

Table 3. **Ablation Study: offsets.** Note that the offsets represent the output of the deformation network. Multi-staged canonical spaces produce better image quality.

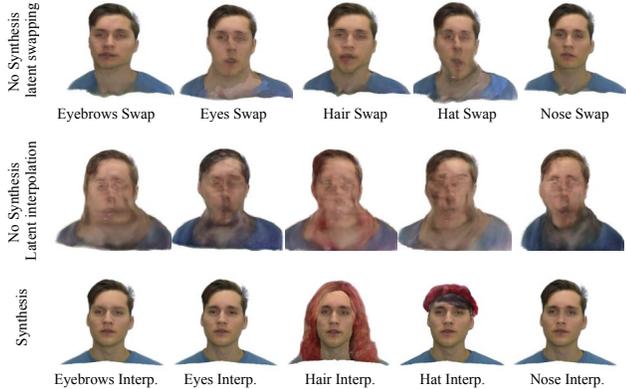


Figure 7. **Qualitative Comparison.** We compare the generative performance of PEGASUS. Our latent code configuration and synthetic DB generation show compelling generative performance in latent interpolation and swapping compared to other baselines.

face part variations while keeping the identity. As expected, both interpolation strategies of the baseline models struggle to generate realistic avatars for the interpolated latent codes.

6.3. Ablation Studies and More Results

Multi-Stage Canonical Space. We compare our multi-stage canonical space framework with the alternative framework with one- or two-stage by PointAvatar frameworks. For quantitative comparison on multi-stage canonical space, we use PSNR, SSIM, and LPIPS [45] metrics. We evaluate them on unseen test sequences with novel head poses and facial expressions from all synthesized videos as shown in Fig. 4. In Tab. 3, our multi-stage canonical space and point deformation outperforms the others stage deformation of all metrics. **More Qualitative Results.** We further demonstrate the performance of our methods by showing the ability to control multiple parts as shown in Fig. 8, and also by showing more interpolation results as shown in Fig. 9 and our Supp. Video.

7. Discussion

We present a method for constructing personalized generative 3D face avatars from monocular video sources. Our compositional generative model enables disentangled controls to selectively alter the facial attributes of the target individual while preserving the identity. Notably, our personalized generative model is built exclusively from monocular videos, without relying on complex multi-view system se-

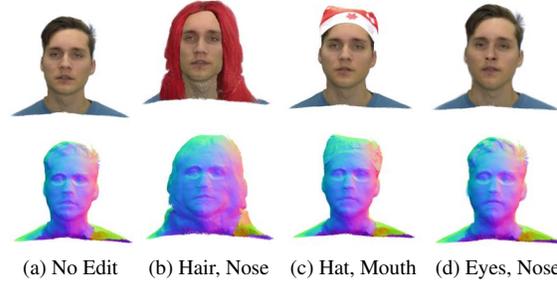


Figure 8. **Multiple Composition.** PEGASUS generates the avatar with multiple face attributes.

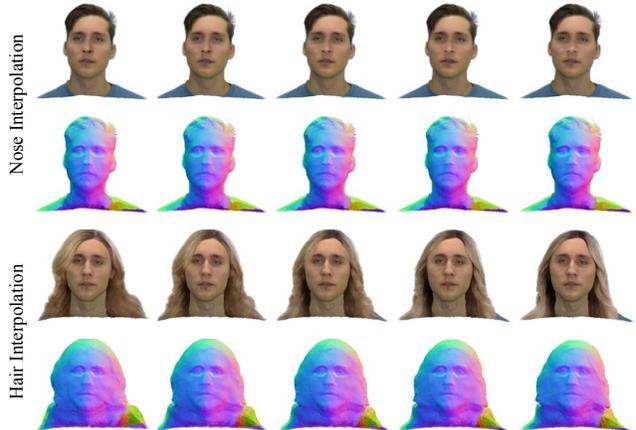


Figure 9. **Person-specific Interpolation.** We interpolate the attribute latent code \mathbf{z}^p , $p \in \{\text{hair, nose}\}$ between two avatars.

tups. To achieve this goal, we present a method to construct a person-specific generative 3D avatar by building a synthetic video collection of the target identity with varying facial attributes, where the videos are synthesized by borrowing parts from diverse individuals from other monocular videos. We also show a zero-shot approach to achieve the same generative modeling more efficiently. For future research, building a more generative model to include multiple identities in a single model can be another exciting extension of our model.

Limitation. As a limitation, the quality of our personalized avatar still does not reach the photo-realistic quality, showing noticeable artifacts. Also, due to the reliance on non-physical-based methods for generating the synthetic DB, our approach exhibits limitations in achieving physical accuracy. We show the failure cases and limitations of the synthetic DB generation in Supp. Mat.

Acknowledgements. We thank Inhee Lee for the fruitful discussion and Hyunwoo Cha for an essential role in collecting and processing in-the-wild videos. This work was supported by SNU Creative-Pioneering Researchers Program, NRF grant funded by the Korean government (MSIT) (No. 2022R1A2C2092724), and IITP grant funded by the Korean government (MSIT) (No.2022-0-00156, No.2021-0-01343). H. Joo is the corresponding author.

References

- [1] Y. Bai, Y. Fan, X. Wang, Y. Zhang, J. Sun, C. Yuan, and Y. Shan. High-fidelity facial avatar reconstruction from monocular video with generative priors. In *CVPR*, 2023.
- [2] S. Bharadwaj, Y. Zheng, O. Hilliges, M. J. Black, and V. Fernandez-Abrevaya. Flare: Fast learning of animatable and relightable mesh avatars. *arXiv preprint arXiv:2310.17519*, 2023.
- [3] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH*, 1999.
- [4] T. Brooks, A. Holynski, and A. A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023.
- [5] C. Cao, Y. Weng, S. Zhou, Y. Tong, and K. Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):413–425, 2013.
- [6] E. R. Chan, C. Z. Lin, M. A. Chan, K. Nagano, B. Pan, S. De Mello, O. Gallo, L. J. Guibas, J. Tremblay, S. Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *CVPR*, 2022.
- [7] S. Chang, G. Kim, and H. Kim. Hairnerf: Geometry-aware image synthesis for hairstyle transfer. In *ICCV*, 2023.
- [8] X. Chen, Y. Zheng, M. J. Black, O. Hilliges, and A. Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *CVPR*, 2021.
- [9] R. Daněček, M. J. Black, and T. Bolkart. Emoca: Emotion driven monocular face capture and animation. In *CVPR*, 2022.
- [10] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- [11] Y. Feng, H. Feng, M. J. Black, and T. Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM TOG*, 40(4):1–13, 2021.
- [12] Y. Feng, J. Yang, M. Pollefeys, M. J. Black, and T. Bolkart. Capturing and animation of body and clothing from monocular video. In *SIGGRAPH ASIA*, 2022.
- [13] Y. Feng, W. Liu, T. Bolkart, J. Yang, M. Pollefeys, and M. J. Black. Learning disentangled avatars with hybrid 3d representations. *arXiv preprint arXiv:2309.06441*, 2023.
- [14] G. Gafni, J. Thies, M. Zollhofer, and M. Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *CVPR*, 2021.
- [15] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *NeurIPS*, 2014.
- [16] P.-W. Grassal, M. Prinzler, T. Leistner, C. Rother, M. Nießner, and J. Thies. Neural head avatars from monocular rgb videos. In *CVPR*, 2022.
- [17] C. Guo, T. Jiang, X. Chen, J. Song, and O. Hilliges. Vid2avatar: 3d avatar reconstruction from videos in the wild via self-supervised scene decomposition. In *CVPR*, 2023.
- [18] A. Haque, M. Tancik, A. A. Efros, A. Holynski, and A. Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. *arXiv preprint arXiv:2303.12789*, 2023.
- [19] A. Hertz, R. Mokady, J. Tenenbaum, K. Aberman, Y. Pritch, and D. Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022.
- [20] H.-I. Ho, L. Xue, J. Song, and O. Hilliges. Learning locally editable virtual humans. In *CVPR*, 2023.
- [21] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33, 2020.
- [22] K. Jiang, S.-Y. Chen, F.-L. Liu, H. Fu, and L. Gao. Nerffaceedit: Disentangled face editing in neural radiance fields. In *SIGGRAPH ASIA*, 2022.
- [23] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [24] H. Kim, G. Lee, Y. Choi, J.-H. Kim, and J.-Y. Zhu. 3d-aware blending with generative nerfs. *arXiv preprint arXiv:2302.06608*, 2023.
- [25] T. Kim, C. Chung, Y. Kim, S. Park, K. Kim, and J. Choo. Style your hair: Latent optimization for pose-invariant hairstyle transfer via local-style-aware hair alignment. In *ECCV*, 2022.
- [26] T. Kim, S. Saito, and H. Joo. Ncho: Unsupervised learning for neural 3d composition of humans and objects. *arXiv preprint arXiv:2305.14345*, 2023.
- [27] N. Kumari, B. Zhang, R. Zhang, E. Shechtman, and J.-Y. Zhu. Multi-concept customization of text-to-image diffusion. In *CVPR*, 2023.
- [28] J. Li, S. Saito, T. Simon, S. Lombardi, H. Li, and J. Saragih. Megane: Morphable eyeglass and avatar network. In *CVPR*, 2023.
- [29] T. Li, T. Bolkart, M. J. Black, H. Li, and J. Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017.
- [30] Z. Liu, M. Li, Y. Zhang, C. Wang, Q. Zhang, J. Wang, and Y. Nie. Fine-grained face swapping via regional gan inversion. In *CVPR*, 2023.
- [31] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021.
- [32] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [33] P. Pérez, M. Gangnet, and A. Blake. Poisson image editing. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 2023.
- [34] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.
- [35] S. Sanyal, T. Bolkart, H. Feng, and M. J. Black. Learning to regress 3d face shape and expression from an image without 3d supervision. In *CVPR*, 2019.
- [36] E. Sella, G. Fiebelman, P. Hedman, and H. Averbuch-Elor. Vox-e: Text-guided voxel editing of 3d objects. In *ICCV*, 2023.
- [37] Y. Shi, X. Yang, Y. Wan, and X. Shen. Semanticstylegan: Learning compositional generative priors for controllable image synthesis and editing. In *CVPR*, 2022.
- [38] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

- [39] Y. Xu, Y. Yin, L. Jiang, Q. Wu, C. Zheng, C. C. Loy, B. Dai, and W. Wu. Transeditor: Transformer-based dual-space gan for highly controllable facial editing. In *CVPR*, 2022.
- [40] Y. Xu, H. Zhang, L. Wang, X. Zhao, H. Huang, G. Qi, and Y. Liu. Latentavatar: Learning latent expression code for expressive neural head avatar. *arXiv preprint arXiv:2305.01190*, 2023.
- [41] T. Yenamandra, A. Tewari, F. Bernard, H.-P. Seidel, M. Elgharib, D. Cremers, and C. Theobalt. i3dmm: Deep implicit 3d morphable model of human heads. In *CVPR*, 2021.
- [42] F. Yin, Y. Zhang, X. Cun, M. Cao, Y. Fan, X. Wang, Q. Bai, B. Wu, J. Wang, and Y. Yang. Styleheat: One-shot high-resolution editable talking face generation via pre-trained stylegan. In *ECCV*, 2022.
- [43] C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, and N. Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *ECCV*, 2018.
- [44] J. Zhang, A. Siarohin, Y. Liu, H. Tang, N. Sebe, and W. Wang. Training and tuning generative neural radiance fields for attribute-conditional 3d-aware face generation. *arXiv preprint arXiv:2208.12550*, 2022.
- [45] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [46] Y. Zheng, V. F. Abrevaya, M. C. Bühler, X. Chen, M. J. Black, and O. Hilliges. Im avatar: Implicit morphable head avatars from videos. In *CVPR*, 2022.
- [47] Y. Zheng, W. Yifan, G. Wetzstein, M. J. Black, and O. Hilliges. Pointavatar: Deformable point-based head avatars from videos. In *CVPR*, 2023.
- [48] J. Zhu, Y. Shen, D. Zhao, and B. Zhou. In-domain gan inversion for real image editing. In *ECCV*, 2020.
- [49] P. Zhu, R. Abdal, J. Femiani, and P. Wonka. Barbershop: Gan-based image compositing using segmentation masks. *arXiv preprint arXiv:2106.01505*, 2021.
- [50] P. Zhu, R. Abdal, J. Femiani, and P. Wonka. Hairnet: Hairstyle transfer with pose changes. In *ECCV*, 2022.
- [51] J. Zhuang, C. Wang, L. Lin, L. Liu, and G. Li. Dreameditor: Text-driven 3d scene editing with neural fields. In *SIGGRAPH ASIA*, 2023.
- [52] W. Zielonka, T. Bolkart, and J. Thies. Towards metrical reconstruction of human faces. In *ECCV*, 2022.