

## My3DGen: A Scalable Personalized 3D Generative Model

Luchao Qi<sup>1</sup> Jiaye Wu<sup>2</sup> Annie N. Wang<sup>1</sup> Shengze Wang<sup>1</sup> Roni Sengupta<sup>1</sup>

<sup>1</sup>The University of North Carolina at Chapel Hill, USA <sup>2</sup>The University of Maryland, USA

{lqi, awang13, shengzew, ronisen}@cs.unc.edu jiyawu@cs.umd.edu

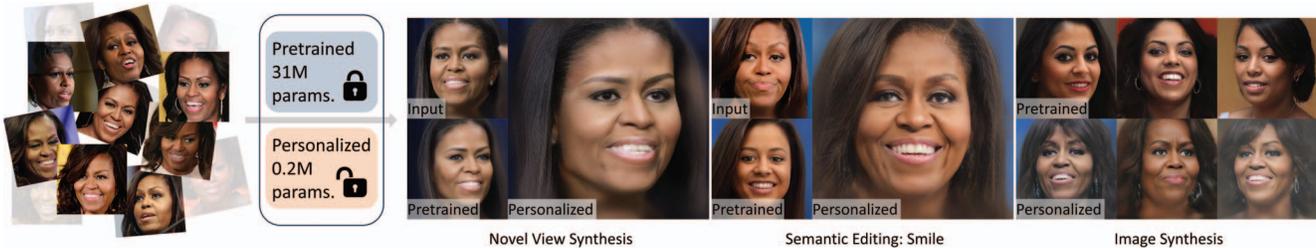


Figure 1. Given 50 images of *Michelle Obama*, we personalize a pre-trained 3D generative prior and demonstrate the benefits of personalization in various downstream tasks. Each downstream task presents the original input image of *Michelle* (top left), alongside the corresponding output generated using the pre-trained face prior (bottom left), compared to the output using our personalized face prior (right). Our personalized prior can faithfully retain the key facial characteristics of *Michelle Obama*, as opposed to the pre-trained prior.

### Abstract

In recent years, generative 3D face models (e.g., EG3D) have been developed to tackle the problem of synthesizing photo-realistic faces. However, these models are often unable to capture facial features unique to each individual, highlighting the importance of personalization. Some prior works have shown promise in personalizing generative face models, but these studies primarily focus on 2D settings. Also, these methods require both fine-tuning and storing a large number of parameters for each user, posing a hindrance to achieving scalable personalization. Another challenge of personalization is the limited number of training images available for each individual, which often leads to overfitting when using full fine-tuning methods. Our proposed approach, My3DGen, generates a personalized 3D prior of an individual using as few as 50 training images. My3DGen allows for novel view synthesis, semantic editing of a given face (e.g. adding a smile), and synthesizing novel appearances, all while preserving the original person’s identity. We decouple the 3D facial features into global features and personalized features by freezing the pre-trained EG3D and training additional personalized weights through low-rank decomposition. As a result, My3DGen introduces only **240K** personalized parameters per individual, leading to a **127×** reduction in trainable parameters compared to the **30.6M** required for fine-tuning the entire parameter space. Despite this significant reduction in storage, our model preserves identity features without compromising the quality of downstream applications, both quantitatively and qualitatively.

### 1. Introduction

Recently, dramatic advancements in deep generative models like generative adversarial networks (GANs) [5, 33, 55, 58] and diffusion models [18, 29] have led to a surge in their popularity for computer vision applications. Notably, GANs have proven to be particularly powerful at generating realistic photos of faces [34–38]. These techniques have been extended to 3D vision as well, leading to the development of models that are capable of reconstructing 3D models of existing 2D facial images, synthesizing novel appearances, and editing various facial attributes such as facial expressions [11, 49, 50, 54, 64]. These models play a crucial role in virtual communication, AR/VR/MR, and content creation, thereby increasing engagement.

However, current 3D generative models [11, 49, 50, 54, 64] are unable to create authentic 3D face models of particular subjects. Although these models can synthesize photo-realistic fake faces, they cannot generate, reconstruct, or modify the distinctive traits of a particular real person’s face without distorting their identity. This problem is further exaggerated for underrepresented demographics whose facial characteristics are sparsely represented in widely used training datasets like FFHQ [37] or CelebA [46]. Existing inversion techniques [1, 6, 39, 60, 69, 73] are only able to preserve identity by tuning the model separately for every test image, an impractical and inefficient approach. Furthermore, they cannot edit the inverted image or synthesize novel appearances without identity distortion. Thus, this paper presents an approach to create a personalized 3D generative prior for an individual. This prior enables 3D facial reconstruction, synthesis of novel appearances, and editing of existing ap-

pearances, while maintaining the individual’s identity.

A major obstacle to the deployment of personalized 3D generative models at scale for real-world applications is their enormous storage demand. We can illustrate this scalability problem with a concrete example: Consider that naively personalizing a 3D generative model such as EG3D [11] requires storing approximately 31 million parameters (121 MB) for each user. For three billion users (monthly active users of Facebook), this would require 363 PB of memory, an extremely cost-prohibitive demand even for a company the size of Meta. Hence, we need to design more parameter-efficient personalization techniques that will enable the building of 3D generative priors for a large population. Another challenge of personalization is that an average user often takes only a limited number of photos of themselves each day or week [48]. The limited training set size makes it difficult to fine-tune a large global generative model, leading to overfitting, mode collapse [2], and data drift [42, 47], all of which prevent the model from generalizing to unseen test images of an individual.

Our idea is to decouple the facial features of an individual into (a) shared global features that can be represented by a generative model trained across many different identities and (b) personalized features of an individual that can be represented with much fewer trainable parameters, trained on images of that individual only. We use EG3D [11] as our pre-trained generative model and train additional weights for low-rank decompositions of every convolutional and fully-connected layer to capture the personalized features of that individual. Drawing inspiration from the recent success of Low-Rank Adaption (LoRA) [31] in parameter-efficient fine-tuning of large language models [32, 76] and diffusion models [63], we use LoRA during personalization, which hasn’t been well explored in convolution-heavy GAN-based models before. Our approach allows for personalization using only 240K (0.9 MB) trainable parameters instead of requiring fine-tuning of all 31 million (121 MB) parameters. For three billion users, this means that only 2.7 PB of storage memory would be needed instead of 363 PB.

Quantitative and qualitative analyses show that our proposed My3DGen outperforms the pre-trained 3D generative model EG3D [11] on multiple tasks including 3D reconstruction, novel appearance synthesis, image enhancement, and semantic editing. Additionally, our personalized model can produce results similar to those achieved by naively fine-tuning a pre-trained model with 31 million parameters [11], while only using as few as 240K trainable parameters. We further provide insights based on different parameterization strategies. We observe that increasing the rank of LoRA modules only contributes to better overfitting of the background without improving the shape or the identity-preserving performance after personalization. Furthermore, we find that personalizing earlier layers of Style-

GAN2 has the most impact on the resulting quality, indicating that coarse and middle layers are more responsible for capturing the shape and identity of the person.

In conclusion, we propose the following contributions. (1) To the best of our knowledge, our research presents the first attempt to personalize a 3D generative model and demonstrates its effectiveness in various downstream tasks; (2) We design a generative prior where an individual’s facial features are disentangled into global features—represented by a pre-trained 3D generative model, and personalized features—captured by training additional low-rank weights requiring only 240K stored parameters. This method helps to avoid overfitting and has the potential to improve performance over naive tuning without LoRA; (3) Departing from previous works that primarily utilized LoRA solely for linear layers within transformers or diffusion models, our approach incorporates LoRA into convolution-centric GAN architectures, presenting an innovative perspective on using LoRA for 3D personalization.

## 2. Related Work

**3D Face Reconstruction.** Parametric 3D models, also known as 3D Morphable Models (3DMM) [8, 22], are often widely used for modeling 3D human faces through a linear combination of basic shapes. Typically, 3DMMs are constructed using high-quality facial scans from multiple individuals. However, when fitting a 3DMM to a test image [41, 67, 68], the result often appears unrealistic due to limitations in the expressiveness of the linear blended model. Also, training a 3DMM using 2D photo collections is challenging due to the absence of corresponding 3D scans. As a result, recent research has focused on developing generative models for 3D faces [12, 25, 54], which are able to achieve a remarkable level of facial detail and capture small wrinkles and bumps by training solely on 2D images. Nevertheless, these models only aim to generate arbitrary fake faces, and there is a lack of research addressing the personalization of such models.

**Personalized Generative Models.** Personalization plays a significant role in the field of generative AI, with diverse applications including deepfakes, video avatars, talking heads, and text-to-image generation [2, 13, 14, 56, 62, 63, 72, 75]. Previous inversion works [3, 6, 60] designed a model-dependent encoder or fine-tuned the pre-trained model to best fit an input image. However, these methods struggle to generate diverse appearances or edit existing appearances of an individual without distorting their identity. Most inversion methods also require updating and storing a separate set of network weights for every single image, which is impractical. In personalizing a 3D generative model, our goal is to use only a single set of network weights that can reconstruct, edit, and synthesize any novel appearances of an individual while preserving identity. Our work is inspired by the recent success in 2D generative model personaliza-

tion [51,78]. However, 2D personalization cannot create 3D faces or produce unique perspectives, and frequently falters for nonfrontal views due to a lack of geometric information. Our paper extends personalization from 2D to 3D, assuming that an individual’s facial appearance can be decomposed into global features represented by the pre-trained model and personalized characteristic features represented by parameter-efficient low-rank adaptive weights.

**Parameter-Efficient Fine-Tuning.** Large foundation models [9, 20, 45, 59, 61] often achieve impressive performance for tasks in their domain. However, the huge number of parameters in such models often prevents them from being fine-tuned for downstream tasks using a limited budget. For example, GPT-4 [53] contains 1.76 trillion parameters, which is impractical for most users to fine-tune. To tackle this issue, many parameter-efficient fine-tuning (PEFT) techniques have been previously proposed to fine-tune models efficiently. In natural language processing, approaches [30, 31, 44] have been proposed to enable efficient adaptation of pre-trained language models to various downstream applications without fine-tuning all of the model’s parameters. For image generation, ControlNet [79], HyperDreamBooth [63], and AnimateDiff [27] have been proposed to tune pretrained diffusion models with additional modules. Our approach is inspired by LoRA [31], a technique for efficiently finetuning large language foundation models by imposing low-rank structures on weight matrices. LoRA [31] was first proposed to predict additional network weights without changing the pre-trained transformer, allowing efficient adaptation and storage for task-specific models [31]. Subsequently, LoRA has found widespread use in fine-tuning pre-trained networks for various downstream tasks, including network quantization, parameter budget allocation, and continual learning [17, 66, 80].

### 3. Method

Our objective is to personalize a 3D generative model with ideally 50 images of an individual. First, we introduce background concepts in Sec. 3.1. Then we formulate the problem of personalizing EG3D [11], a pre-trained 3D generative model, in Sec. 3.2. Finally, we discuss how we incorporate parameter-efficient personalization into our model in Section 3.3 with a visual overview in Fig. 2.

#### 3.1. Preliminaries

**EG3D.** In this work we use the state-of-the-art 3D generative model EG3D [11] as the pre-trained model that captures global facial features across multiple identities. EG3D has four main components: (i) a StyleGAN2 generator that takes as input a random latent code and outputs  $256 \times 256 \times 96$  feature maps, which are further reshaped into three  $256 \times 256 \times 32$  triplanes, (ii) a neural renderer that decodes triplane features and renders a face given a camera pose input, (iii) a super-resolution module that up-samples rendered images to a resolution of  $512 \times 512$ , and

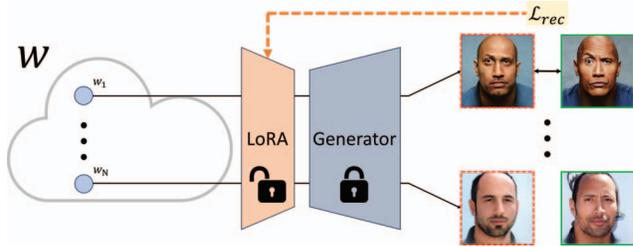


Figure 2. Architecture of our personalization approach. We project an individual’s images into StyleGAN2’s  $\mathcal{W}$  space through latent code optimization to obtain a set of latent anchors. We then tune the generator to reconstruct an individual’s images based on these anchors. During tuning, the generator is frozen while only LoRA weights are personalized.

(iv) a StyleGAN2 discriminator that differentiates between generated images and real images. The entire pipeline is trained following the typical non-saturating minimax GAN loss [24]. In this paper, we discard the discriminator and will use a reconstruction loss to personalize the generator.

**LoRA.** Our approach is based on Low-Rank Adaptation of Large Language Models (LoRA) [31], a technique for efficiently finetuning convolution-free transformer networks by imposing low-rank structures on weight matrices. LoRA shows that a pre-trained model’s weight matrix  $W_0 \in \mathbb{R}^{d \times k}$  for any fully-connected layer can be fine-tuned with a low-rank decomposition using the following formulation: Let  $W_{ft} = W_0 + \Delta W = W_0 + BA$  be the fine-tuned adapted weight matrix, where  $B \in \mathbb{R}^{d \times r}$  and  $A \in \mathbb{R}^{r \times k}$  are trained while  $W_0$  is frozen, and the rank  $r \ll \min(d, k)$  [31]. Although this approach demonstrates impressive decomposition performance for transformers [28] and diffusion models [63], previous work solely focuses on the decomposition of linear layers [27, 40, 63, 66, 71], while GAN-based models are convolution-heavy. In our paper, we apply LoRA to the convolution-heavy EG3D model for personalization.

#### 3.2. Personalization Formulation

We consider a scenario where an individual has a training set of  $N$  2D images  $x_i$ , each with an associated camera pose  $c_i$ , denoted as  $\mathcal{D}_p = \{(x_i, c_i)\}_{i=1}^N$ . We let  $G(\cdot; \theta_G + \Delta\theta_G)$  denote our personalized EG3D model where  $G$  is the EG3D architecture,  $\theta_G$  represents the frozen pretrained weights, and  $\Delta\theta_G$  represents our trainable parameters. Our goal is to tune this model on a low-dimensional manifold in the  $\mathcal{W}$  latent space, dubbed a personal convex hull [51], which is a subspace defined by all latent codes of the  $N$  images. Our personalization scheme is inspired by the success of PTI [60] and MyStyle [51] in tuning a pretrained 2D StyleGAN. We first invert images from  $(x_i, c_i) \sim \mathcal{D}_p$  with associated camera poses into latent vectors  $w_i$ , dubbed anchors in MyStyle [51], with an off-the-shelf inversion technique [38]. That is, we freeze both the camera pose and the model weights while only optimizing the randomly initialized latent vector.

We then tune the weights  $\Delta\theta_G$  of the model  $G$  using reconstruction objective  $\mathcal{L}_{rec}$ , which represents the difference between each image  $x_i$  and its reconstruction. Formally,

$$\mathcal{L}_{rec}^{(i)} = \mathcal{L}_{l_{lips}}(G(w_i, c_i; \theta_G + \Delta\theta_G), x_i) + \lambda_{\mathcal{L}_2} \|G(w_i, c_i; \theta_G + \Delta\theta_G) - x_i\|_2 \quad (1)$$

and we have  $\Delta\theta_{G_p} = \operatorname{argmin}_{\Delta\theta_G} \sum_{i=1}^N \mathcal{L}_{rec}^{(i)}$  as our optimized parameters for the personalized EG3D model  $G_p$ .

### 3.3. Parameter Efficient Personalization

Naively fine-tuning the model can lead to forgetting of knowledge learned in pretraining and entanglement of pre-trained facial features with the personalized features, compromising both the model’s interpretability and generalization. Therefore, in order to decouple the global features, captured in the pre-trained model, from the personalized features, we freeze the weights of the pre-trained model and tune additional weights to capture an individual’s distinct facial priors. We use the technique of Low-Rank Adaptation (LoRA) [31] to train the additional personalized weights, using only 240K parameters per identity.

The original LoRA paper focuses on convolution-free transformers where linear fully connected layers are decomposed into low-rank submatrices  $A$  and  $B$ , respectively [31]. However, EG3D’s generator StyleGAN2 and super-resolution modules are convolution-heavy, and convolution operations are often implemented with general matrix multiplication using the well-known ‘im2col’ algorithm [4]. Thus, when adapting for personalization, we decompose both the convolution and fully connected layers in the StyleGAN2 generator and super-resolution modules using LoRA<sup>1</sup>. Pretrained weights  $W_0$  of the StyleGAN2 generator and the super-resolution module are frozen while only  $A$  and  $B$  are trained. We tune all the parameters of the neural renderer in EG3D, as it is relatively small with 4K parameters. Under this setting, the number of parameters needed to be tuned is determined by the rank  $r$ . With rank  $r = 1$ , we can reduce the 30.6M trainable parameters of EG3D to only 240k, a reduction of  $127\times$ .

## 4. Experiments

We first discuss the details of our evaluation framework including datasets, experiment pipeline, and evaluation metrics in Sec. 4.1. Then we present both qualitative and quantitative results of our approach, My3DGen, in Sec. 4.2. We further analyze the effect of LoRA for personalization in Sec. 4.3. Finally, we present ablation studies in Sec. 4.4.

### 4.1. Experiment Details

**Dataset.** We conduct experiments using facial images of celebrities, the same dataset used by Mystyle [51]. Images are preprocessed following [11, 34] to align and crop faces

<sup>1</sup>Details of our implementation of the convolutional LoRA decomposition are provided in the supplemental material.

to  $512 \times 512$ . Finally, the faces are separated into reference sets and test sets for each celebrity. The number of images included in the reference set and the test set for each celebrity is presented in the supplementary. Unless otherwise noted, for each celebrity, we personalize a model using 50 images from the reference set as the training set.

We use an off-the-shelf pose extraction [16] model to both identify the face region and label the pose of the face, following the same pipeline in EG3D [11]. We project images into the model’s  $\mathcal{W}$  latent space following the StyleGAN [38] optimization scheme for 500 iterations to obtain their latent codes (anchors).

**Training.** We choose 50 images per individual for all personalization tunings. The effect of dataset size is further discussed in Sec. 4.4. We tune the model starting from the pre-trained EG3D (on FFHQ), which outputs images at resolutions of both  $128 \times 128$  ( $I_{128}$ ) and  $512 \times 512$  ( $I_{512}$ ). Hyperparameters are chosen as follows:  $\lambda_{lips} = \lambda_{\mathcal{L}_2} = 1$ . We apply  $\mathcal{L}_{rec}$  on both  $I_{128}$  and  $I_{512}$ . Our further experiments show that a higher LPIPS regularization weight ( $\lambda_{lips} > 1$ ) will lead to checkerboard-style artifacts and a lower LPIPS weight ( $\lambda_{lips} < 0.1$ ) will cause nonphotorealistic artifacts, which aligns with results in previous work [6]. Based on this observation, we use  $\lambda_{lips} = 1$ .

**Evaluations.** We evaluate our methods on unseen test images for the following primary downstream tasks: 1) *Image inversion* through PTI [60]; 2) *Image interpolation* where we interpolate between two anchors in the latent space and generate images that morph from one face to another; 3) *Image synthesis* where the goal is to generate novel appearances of an individual by sampling from a latent space, following the protocol outlined in Mystyle [51]; 4) *Image enhancement* tasks such as image-inpainting and super-resolution; 5) *Semantic editing* where the goal is to modify the facial expression or age of the person while maintaining the identity and pose. Note that we modify the model weights after personalization via PTI only for image-inversion tasks, for consistency with prior research on 3D GAN [11]. For any novel views, we render faces with a yaw range of  $\pm 0.35$  (radians) and a pitch range of  $\pm 0.25$  (radians) relative to the frontal face for all our experiments.

**Metrics.** When evaluating inversion outcomes, it is standard to use pixel-based metrics such as PSNR and SSIM or neural network-based perceptual metrics such as LPIPS [81] and DISTS [19] to assess the inverted image and compare it with the original image. However, due to errors in the estimation of face poses, Live3Dportrait [69] reports that a minor misalignment between the ground truth and the estimated face poses will cause traditional pixel-based metrics to be unreliable. Recent work further indicates that deep perceptual image metrics are also sensitive to small misalignments [23]. Therefore, we additionally adopt the facial identity score  $ID_{sim}$ , *i.e.* a metric to evaluate the preserva-

	Input	Output (input view)	Output (novel views)		
Pretrained					
Full F.T.					
Ours					
Pretrained					
Full F.T.					
Ours					

F.T.	LoRA	DISTS ↓	LPIPS ↓	ID <sub>sim</sub> ↑	Personal. Params. ↓
-	-	0.12	0.19	0.56	-
✓	-	0.08	0.12	0.60	31M
✓	✓	0.08	0.13	0.61	0.2M

Figure 3. Qualitative (top) and quantitative (bottom) evaluation for image inversion, *i.e.* generate 3D-aware view synthesis from a single image. F.T. indicates fine-tuning, and Personal. Params. indicates the number of personalized parameters. DISTS and LPIPS compare the inverted image against the input image in the same pose, ID<sub>sim</sub> evaluates the preservation of identity across multiple poses. Ours is My3DGen with LoRA rank  $r = 1$  where the number of trainable parameters = 0.2M. Visual differences are highlighted with a red-box, zoom in to view finer details.

tion of the individual’s identity [10, 51]. Nonetheless, we still include DISTS and LPIPS results in our inversion tasks to align with prior works [11, 69]. The identity score ID<sub>sim</sub> of an image is determined by the individual’s reference set, which contains all the images available from that individual. Given an image, we extract the identity features and report the cosine similarity of the given image to the nearest image in the personal reference set. Formally,

$$\text{ID}_{sim}(w_i, c_i) = \max \{ \langle R(G(w_i, c_i, \cdot)), R(x_j) \rangle \}_{j=1}^N \quad (2)$$

where  $R$  is a pretrained ArcFace [15] network for feature recognition,  $\langle \cdot, \cdot \rangle$  computes the cosine similarity between its argument as the ID scores [57], and  $N$  is the number of 2D images in a reference set.

## 4.2. Applications of My3DGen

**Inversion.** Given a 2D RGB image of a face, image inversion here refers to performing 3D reconstruction of the face. We apply PTI [60] for inversion benchmarking to align with the original EG3D work [11]. We optimize a randomly initialized latent code for 600 iterations, followed by fine-tuning the model for an additional 350 iterations [69]. We calculate DISTS and LPIPS for single-view inversion, and report ID<sub>sim</sub> to evaluate multi-view reconstruction.

As shown in Fig. 3, the pretrained EG3D can have artifacts such as i) a visible seam between the face and the rest of the head (1st row, 3rd column), ii) 3D shape distortion (4th row, 5th column), and iii) identity drift caused by changing pose, which can be observed in Fig. 1 regarding novel view synthesis. Furthermore, the pretrained EG3D overly smooths skin textures. We recommend readers to zoom in to observe these finer details (2nd column of Fig. 3). In contrast, ours decomposes facial features into global features and personalized features, producing a higher quality and identity-preserving inversion. Compared to full fine-tuning, our approach yields almost the same quality for single-view inversion as measured by LPIPS, and produces similar multiview inversion results with a slightly higher ID<sub>sim</sub>. We speculate that full fine-tuning can result in better single-view quality by inverting images with an overfitted background, but may also introduce artifacts by overfitting particular facial features. For background overfitting, we examine this problem in Fig. 8 where we demonstrate that increasing LoRA rank enhances background fitting but provides minor improvements in terms of ID<sub>sim</sub>. For facial features, as shown in Fig. 3 (5th row), full fine-tuning does not preserve the identity in the lower jaw and ear regions, which are obscured in the input image. My3DGen can help reduce these artifacts as well as ‘floaters’ from NeRF [74]. We suggest the readers check the video results in the supplementary for more examples.

**Interpolation.** The results obtained from the inversion tasks are overfitted to the test image via PTI, requiring separate optimization of the generator for each image. Our approach aims to provide a single generative model for all images of an individual without requiring further optimization. We support this claim by generating images from latent vectors produced through linear interpolation between two randomly selected training anchors [51].

We interpolate between the anchor pair at 10 equally spaced interpolation weights. At each interpolation step  $\theta$ , given the interpolated latent code, we randomly generate 20 novel views and compute the average ID<sub>sim</sub> scores of each view as described in Eq. 2. We repeat this process for 10 randomly selected anchor pairs for each personalized model, averaging our results across all models and pairs. We compare the results before and after personalization in Fig. 4 and show that personalization maintains the identity

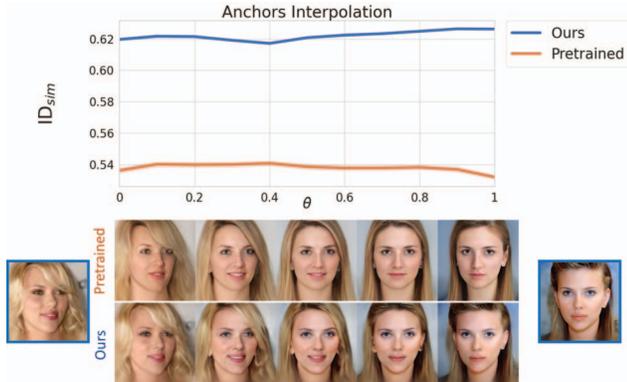


Figure 4. Qualitative (top) and quantitative (bottom) evaluation for interpolation in latent space between two anchor images, highlighted in **color**. We measure identity preservation using  $ID_{sim}$ , for which illustrative visual results at different interpolation steps  $\theta$  are also provided. We compare the interpolation results **before** and **after** personalization.

when traversing between two anchor images (exemplified by *Scarlett Johansson*) while the pre-trained model distorts identity.

**Synthesis.** We conduct image synthesis to further examine the personalization capacity of our approach. Our goal is to produce new, distinctive images of an individual that have not been seen before. We sample a latent code from the convex hull following MyStyle [51]. Then we feed the latent code into the generator together with a random pose. To this end, we randomly synthesize images from each generator and compare the results based on identity preservation. We evaluate the image synthesis results and present them in Fig. 5, accompanied by visual examples. To assess the diversity of the synthesized images, we follow the protocol proposed by Ojha et al. [52]. Specifically, for each celebrity, we produce 1,000 images and assign each of them to one of the  $k$  training images, by choosing the one with the lowest LPIPS distance. We then compute the standard deviation of pairwise LPIPS distances within members of the same cluster and then average over the  $k$  clusters. Our model surpasses pretrained EG3D in all metrics and performs comparably to full fine-tuning with fewer trainable parameters.

**Image Enhancement.** We choose the tasks of image inpainting (IP) and super-resolution (SR) as representative examples of image enhancement. To perform inpainting, we utilize a mask positioned at the center of the faces and feed the masked images into the generator. We subsequently post-process the resulting outputs by blending the generator’s output within the masked area with the input in other areas. The background regions of the outputs are replaced through post-processing using a Lanczos-upsampled version of the input image. These regions have been segmented according to the methodology in [70]. For SR, we reduce the original image size from  $512 \times 512$  to  $32 \times 32$  and supply the blurred images to the generator. It is noteworthy that

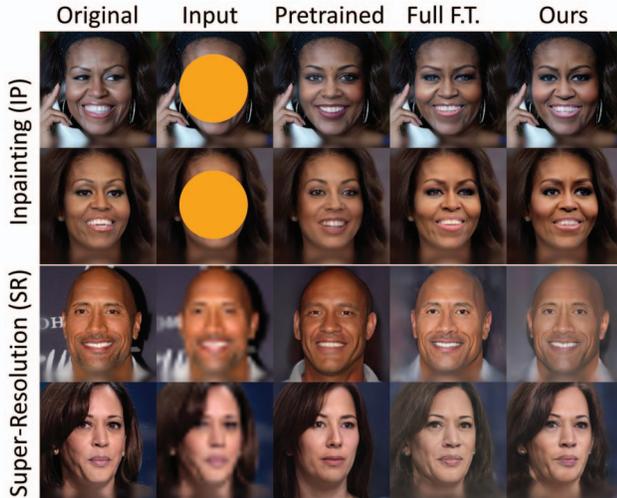


F.T.	LoRA	$ID_{sim} \uparrow$	Diversity $\uparrow$	Personal. Params. $\downarrow$
-	-	0.53	0.19	-
✓	-	0.62	0.21	31M
✓	✓	0.62	0.21	0.2M

Figure 5. Qualitative (top) and quantitative (bottom) evaluation for synthesizing novel appearances of an individual. Generated images of *Oprah Winfrey* are provided for visual inspection. F.T. indicates fine-tuning and Personal. Params. indicates the number of personalized parameters.

in previous studies on image enhancement [82], no quality evaluation is performed against a reference standard, *i.e.* a ground truth image. “Quality” here refers to a pixel-to-pixel comparison against the original ground truth. The reasoning behind this approach is that, although the restored facial details may differ from the original image, personalization still preserves the key facial characteristics of the individual, resulting in valid restorations [51, 78]. We follow the established convention and do not include additional quantitative quality evaluations. However, in addition to using the  $ID_{sim}$  metric in the experiments, we also conducted user studies based on subjective assessments of identity preservation.

We used Amazon Mechanical Turk to gather 330 responses from 17 users. In the study, users were shown the original image, along with two results: one from the pre-trained model and the other from the personalized model, in a randomized order. Users were instructed to select the result that most closely resembled the person in the original image. In cases where both results were similar, participants could select ‘No Preference’. Results are reported as a percentage of each selected option. Images utilized in the user study were selected from a random subset of those used for quantitative evaluation. We present the enhancement results in Fig. 6, showing visual examples of inpainting (IP) and super-resolution (SR) tasks followed by both quantitative and qualitative analysis. The User % for ‘No Preference’ is 5.5% for inpainting and 3.6% for super-resolution. Our results demonstrate that My3DGen successfully personalizes the pretrained EG3D, achieving both a higher quantitative score ( $ID_{sim}$ ) and qualitative preference (User %).

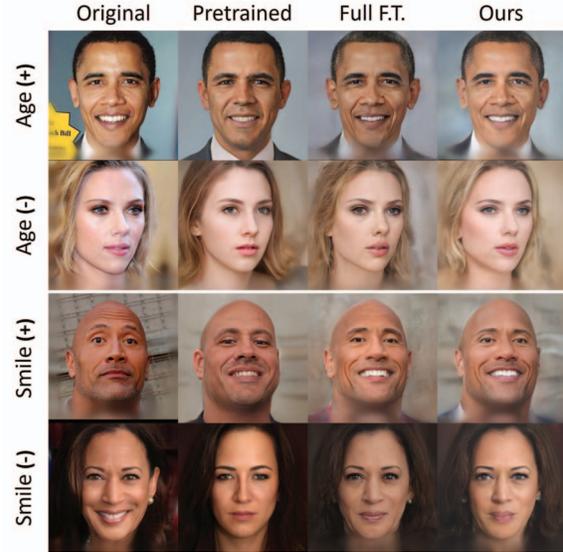


	F.T.	LoRA	ID <sub>sim</sub> ↑	User % ↑	Personal. Params. ↓
IP	-	-	0.62	12.7	-
	✓	-	0.72	-	31M
	✓	✓	0.72	81.8	0.2M
SR	-	-	0.61	7.9	-
	✓	-	0.73	-	31M
	✓	✓	0.73	88.5	0.2M

Figure 6. Qualitative (top) and quantitative evaluation (bottom) for image enhancement by inpainting (IP) and super-resolution (SR). The original images are degraded as input images and then fed into the model. We have included original images for the benefit of readers who may not be familiar with the faces of *Michelle Obama*, *Dwayne Johnson*, and *Kamala Harris*. F.T. indicates fine-tuning and Personal. Params. indicates the number of personalized parameters. User % reflects the percentages of responses for each option in the image enhancement tasks.

**Semantic Editing.** Personalization also enables tailored editing of facial attributes. For example, while a generalized pre-trained model may learn how an average human smiles or ages with time, it struggles to capture the specific smile or aging process of an individual. Our hypothesis is that a personalized generative prior can enable us to more accurately depict how specific attributes of an individual change.

To validate this hypothesis, we perform semantic editing of ‘smile’ and ‘age’ features using 3D generative priors, both pretrained and personalized. We start by identifying suitable editing directions for each model using the InterFaceGAN framework [65], followed by reconstructing the input image and then conducting personalized semantic editing in the  $\alpha$ -space, as described in MyStyle [51]. Our findings are illustrated in Fig. 7, where we also employ the user study as quantitative metrics for evaluation. Our study design for the user survey aligns with that of the image enhancement evaluations, including 165 responses from a total of 17 participants. Quantitatively, our model



	F.T.	LoRA	ID <sub>sim</sub> ↑	User % ↑	Personal. Params. ↓
Age (+)	-	-	0.60	4.8	-
	✓	-	0.76	-	31M
	✓	✓	0.76	95.2	0.2M

Figure 7. Qualitative (top) and quantitative evaluation (bottom) for semantic editing of ‘smile’ and ‘age’. We show *Barack Obama*, *Scarlett Johansson*, *Dwayne Johnson*, and *Kamala Harris* top-to-bottom. F.T. indicates fine-tuning and Personal. Params. indicates the number of personalized parameters. User % reflects the percentages of responses for each option in the downstream task.

yields a nearly identical ID<sub>sim</sub> to the fully fine-tuned EG3D model and substantially higher ID<sub>sim</sub> than the pretrained model. Qualitatively, we note that our model can sometimes outperform the fully fine-tuned model in retaining image styles, such as lighting, background color, and hairstyle post-editing. A detailed examination of Fig. 7 demonstrates improved preservation of skin tone and background color in the first row, enhanced reproduction of hair color, style, and background color in the second row, and better handling of lighting and skin tone in the last row, showing LoRA may mitigate overfitting in semantic editing.

### 4.3. Analysis of Personalization with LoRA

In Sec. 4.2, we show that personalizing 3D-GAN with LoRA can help avoid overfitting and has the potential to improve downstream results such as inversion and semantic editing compared to naive full fine-tuning. In this section, we start by demonstrating that tuning the pre-trained model with LoRA rank  $r = 1$  suffices for personalization. We then study the importance of feature blocks with different resolutions for personalization with rank 1.

**Effect of Rank in LoRA.** To investigate the impact of rank selection in LoRA and find the optimal rank, we repeat the same personalization method for each celebrity, with LoRA

Table 1. Quantitative analysis of the rank of LoRA for inversion and interpolation tasks. ‘-’ indicates full fine-tuning. We report the average  $ID_{sim}$  across latent paths from the interpolation task.

LoRA	Inversion		Interpolation	Personal.
	LPIPS ↓	$ID_{sim}$ ↑	$ID_{sim}$ ↑	Params. ↓
-	0.12	0.60	0.62	31M
rank=1	0.13	0.60	0.62	0.2M
rank=4	0.11	0.60	0.63	0.9M
rank=16	0.09	0.60	0.63	3.5M



	rank: 1	rank: 4	rank: 16
Input	LPIPS: 0.18	LPIPS: 0.15	LPIPS: 0.05
	$ID_{sim}$ : 0.61	$ID_{sim}$ : 0.60	$ID_{sim}$ : 0.61

Figure 8. Visual examples of inversion results for different LoRA ranks. The backgrounds are highlighted to aid inspection along with LPIPS and  $ID_{sim}$  scores.

ranks of 1, 4, and 16, tuning the generator accordingly. We report the model’s performance on inversion and interpolation tasks in Table 1. For inversion tasks, increasing the rank from 1 to 16 doesn’t help improve identity preservation measured by  $ID_{sim}$ , but it does help to fit the background and thus leads to a better LPIPS score, for which visual results are shown in Fig. 8. This clarifies the contrast between our method and full fine-tuning regarding single-view inversion results outlined in the inversion task in Sec. 4.2. However, since matching backgrounds is not the primary focus of personalization and there is only a slight improvement in  $ID_{sim}$  for interpolation tasks when the rank increases. Thus, we choose to use rank = 1 in most of our experiments.

**Personalization of Feature Blocks.** We show that feature blocks with different resolutions have different levels of importance during personalization. Determining the importance of weights in a network is an open question [7]. Previous work proposes including gradient information to evaluate the importance of LoRA modules during the tuning [80]. Here, after tuning with LoRA, we follow the common approach in model pruning where the change in parameter magnitude is used as a criterion to evaluate layer importance [21, 43]. That is,  $|\Delta W|/|W_0| \times 100\%$ , where  $\Delta W$  is the personalized LoRA weights and  $W_0$  is the pre-trained weights. EG3D uses the first seven StyleGAN2 resolution blocks to generate  $256 \times 256$  tri-plane representations. We compute the relative LoRA weight change for each block against the pre-trained weights.

We report the mean and variance of the changes across different individuals in Fig. 9. The results show that feature blocks of resolution  $8 \times 8$ ,  $16 \times 16$ , and  $32 \times 32$  require more weight changes than other resolution blocks. This suggests that distinct facial features are primarily learned in the ‘coarse’ and ‘middle’ layers during personalization, with minimal fine-tuning needed for the ‘fine’ layers.

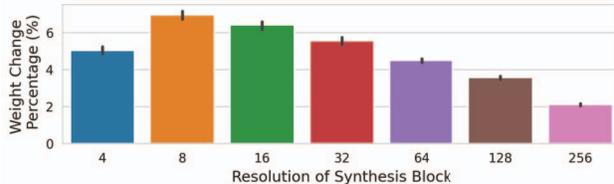


Figure 9. LoRA weight change compared to the pre-trained weights after personalization, averaged across celebrities.

Dataset $\mathcal{D}_p$ Size Ablations	Metric	$\mathcal{D}_p = 10$	$\mathcal{D}_p = 50$	$\mathcal{D}_p = 100$
		$ID_{sim}$ ↑	0.618	0.628

Table 2. The effect of training dataset size  $\mathcal{D}_p$  on personalization. We evaluate each model’s performance through the interpolation task in Sec. 4.2 and report the average  $ID_{sim}$ .

#### 4.4. Ablation study

**Effect of Dataset Size.** We investigate the effect of training dataset size on the performance of the personalized generator. We sample various random subsets of images for each celebrity, with sizes of 10, 50, and 100, and tune the generator accordingly. Next, we assess the generator’s performance using the interpolation tasks outlined in Section 4.2, which evaluates the generalization capacity of the personal convex hull by traversing through the latent subspace. We report the average  $ID_{sim}$  across latent paths. As shown in Table 2, the performance improves significantly from 10 to 50 images, but there is no significant improvement from 50 to 100 images. While further experiments are required, we speculate that adding more images may not contribute to dataset diversity and might even hurt the results [51]. Therefore, we use 50 images for our personalization method in most of our experiments, unless otherwise stated.

**AFHQv2 Cats.** In addition to human faces, we also extend our personalization method to cat faces with the same personalization procedure. We present our results in the supplementary material, highlighting the versatility and effectiveness of our method across different domains.

### 5. Discussion

**Limitations and Future Work.** Our model can accurately capture facial features but encounters difficulty when objects heavily obscure the face (e.g. hats, phones, etc). Further works may leverage the power of per-image inversion methods for better performance [6, 26, 77]. Ours also struggles with heavily cropped faces, where the whole face is not fully captured in the original image and thus image boundaries are filled with reflection-padded values to align the face during preprocessing. One may mask out these invalid values during inversion. This will be similar to the in-painting task discussed in Section 4.2.

**Ethical Considerations.** This study holds the potential to generate manipulated images of actual individuals, posing a substantial societal threat. Future research for detecting fake composites is needed.

## References

- [1] Rameen Abdal, Hsin-Ying Lee, Peihao Zhu, Menglei Chai, Aliaksandr Siarohin, Peter Wonka, and Sergey Tulyakov. 3davatargan: Bridging domains for personalized editable avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4552–4562, 2023. [1](#)
- [2] Mehran Aghabozorgi, Shichong Peng, and Ke Li. Adaptive IMLE for Few-shot Pretraining-free Generative Modelling. 2023. [2](#)
- [3] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Restyle: A residual-based stylegan encoder via iterative refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6711–6720, 2021. [2](#)
- [4] Andrew Anderson, Aravind Vasudevan, Cormac Keane, and David Gregg. High-Performance Low-Memory Lowering: GEMM-based Algorithms for DNN Convolution. In *2020 IEEE 32nd International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, pages 99–106, Sept. 2020. ISSN: 2643-3001. [4](#)
- [5] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223. PMLR, 06–11 Aug 2017. [1](#)
- [6] Ananta R. Bhattarai, Matthias Nießner, and Artem Sevastopolsky. Triplanenet: An encoder for eg3d inversion. *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2024. [1](#), [2](#), [4](#), [8](#)
- [7] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Guttag. What is the State of Neural Network Pruning?, Mar. 2020. arXiv:2003.03033 [cs, stat]. [8](#)
- [8] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, page 187–194, USA, 1999. ACM Press/Addison-Wesley Publishing Co. [2](#)
- [9] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. [3](#)
- [10] Kaidi Cao, Yu Rong, Cheng Li, Xiaoou Tang, and Chen Change Loy. Pose-Robust Face Recognition via Deep Residual Equivariant Mapping. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5187–5196, Salt Lake City, UT, USA, June 2018. IEEE. [5](#)
- [11] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient Geometry-aware 3D Generative Adversarial Networks. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16102–16112, New Orleans, LA, USA, June 2022. IEEE. [1](#), [2](#), [3](#), [4](#), [5](#)
- [12] Eric R. Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5795–5805, Nashville, TN, USA, June 2021. IEEE. [2](#)
- [13] Kun Cheng, Xiaodong Cun, Yong Zhang, Menghan Xia, Fei Yin, Mingrui Zhu, Xuan Wang, Jue Wang, and Nannan Wang. Videoretalking: Audio-based lip synchronization for talking head video editing in the wild. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. [2](#)
- [14] Jun Myeong Choi, Max Christman, and Roni Sengupta. Personalized Video Relighting With an At-Home Light Stage, Dec. 2023. arXiv:2311.08843 [cs]. [2](#)
- [15] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. [5](#)
- [16] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. [4](#)
- [17] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. QLoRA: Efficient Finetuning of Quantized LLMs, May 2023. arXiv:2305.14314 [cs]. [3](#)
- [18] Prafulla Dhariwal and Alex Nichol. Diffusion Models Beat GANs on Image Synthesis, June 2021. arXiv:2105.05233 [cs, stat]. [1](#)
- [19] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P. Simoncelli. Image Quality Assessment: Unifying Structure and Texture Similarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5):2567–2581, May 2022. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. [4](#)
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, June 2021. arXiv:2010.11929 [cs]. [3](#)
- [21] Bryn Elesedy, Varun Kanade, and Yee Whye Teh. Lottery Tickets in Linear Models: An Analysis of Iterative Magnitude Pruning, July 2021. arXiv:2007.08243 [cs, stat]. [8](#)
- [22] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Trans. Graph.*, 40(4), jul 2021. [2](#)

- [23] Abhijay Ghildyal and Feng Liu. Shift-tolerant perceptual similarity metric. In *European Conference on Computer Vision*, pages 91–107. Springer, 2022. 4
- [24] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 3
- [25] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleNeRF: A Style-based 3D-Aware Generator for High-resolution Image Synthesis, Oct. 2021. arXiv:2110.08985 [cs, stat]. 2
- [26] Yuming Gu, Hongyi Xu, You Xie, Guoxian Song, Yichun Shi, Di Chang, Jing Yang, and Linjie Luo. Diffportrait3d: Controllable diffusion for zero-shot portrait view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10456–10465, 2024. 8
- [27] Yuwei Guo, Ceyuan Yang, Anyi Rao, Yaohui Wang, Yu Qiao, Dahua Lin, and Bo Dai. AnimateDiff: Animate Your Personalized Text-to-Image Diffusion Models without Specific Tuning, July 2023. arXiv:2307.04725 [cs]. 3
- [28] Xuehai He, Chunyuan Li, Pengchuan Zhang, Jianwei Yang, and Xin Eric Wang. Parameter-efficient model adaptation for vision transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 817–825, 2023. 3
- [29] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models, Dec. 2020. arXiv:2006.11239 [cs, stat]. 1
- [30] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019. 3
- [31] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models, Oct. 2021. arXiv:2106.09685 [cs] version: 2. 2, 3, 4
- [32] Zhiqiang Hu, Lei Wang, Yihui Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya Poria, and Roy Ka-Wei Lee. LLM-Adapters: An Adapter Family for Parameter-Efficient Fine-Tuning of Large Language Models, Oct. 2023. arXiv:2304.01933 [cs]. 2
- [33] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. 1
- [34] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive Growing of GANs for Improved Quality, Stability, and Variation, Feb. 2018. arXiv:1710.10196 [cs, stat]. 1, 4
- [35] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33:12104–12114, 2020. 1
- [36] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. 1
- [37] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 1
- [38] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 1, 3, 4
- [39] Jaehoon Ko, Kyusun Cho, Daewon Choi, Kwangrok Ryoo, and Seungryong Kim. 3d gan inversion with pose optimization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2967–2976, 2023. 1
- [40] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1931–1941, 2023. 3
- [41] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos Zafeiriou. AvatarMe++: Facial Shape and BRDF Inference With Photorealistic Rendering-Aware GANs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9269–9284, Dec. 2022. Publisher: IEEE Computer Society. 2
- [42] Jason Lee, Kyunghyun Cho, and Douwe Kiela. Countering Language Drift via Visual Grounding, Sept. 2019. arXiv:1909.04499 [cs]. 2
- [43] Jaeho Lee, Sejun Park, Sangwoo Mo, Sungsoo Ahn, and Jinwoo Shin. Layer-adaptive sparsity for the Magnitude-based Pruning, May 2021. arXiv:2010.07611 [cs]. 8
- [44] Zhaojiang Lin, Andrea Madotto, and Pascale Fung. Exploring versatile generative language model via parameter-efficient transfer learning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 441–459, Online, Nov. 2020. Association for Computational Linguistics. 3
- [45] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019. 3
- [46] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. 1
- [47] Yuchen Lu, Soumye Singhal, Florian Strub, Aaron Courville, and Olivier Pietquin. Countering Language Drift with Seeded Iterated Learning. In *Proceedings of the 37th International Conference on Machine Learning*, pages 6437–6447. PMLR, Nov. 2020. ISSN: 2640-3498. 2
- [48] Andrew D. Miller and W. Keith Edwards. Give and take: a study of consumer photo-sharing culture and practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 347–356, San Jose California USA, Apr. 2007. ACM. 2

- [49] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7588–7597, 2019. [1](#)
- [50] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing Scenes as Compositional Generative Neural Feature Fields. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11448–11459, Nashville, TN, USA, June 2021. IEEE. [1](#)
- [51] Yotam Nitzan, Kfir Aberman, Qiurui He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. Mystyle: A personalized generative prior. *ACM Transactions on Graphics (TOG)*, 41(6):1–10, 2022. [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [52] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot image generation via cross-domain correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10743–10752, 2021. [6](#)
- [53] OpenAI. GPT-4 Technical Report, Mar. 2023. arXiv:2303.08774 [cs]. [3](#)
- [54] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13503–13513, 2022. [1](#), [2](#)
- [55] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [1](#)
- [56] Akshay Paruchuri, Xin Liu, Yulu Pan, Shwetak Patel, Daniel McDuff, and Soumyadip Sengupta. Motion Matters: Neural Motion Transfer for Better Camera Physiological Measurement, Nov. 2023. arXiv:2303.12059 [cs]. [2](#)
- [57] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2085–2094, 2021. [5](#)
- [58] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks, Jan. 2016. arXiv:1511.06434 [cs]. [1](#)
- [59] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. [3](#)
- [60] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on graphics (TOG)*, 42(1):1–13, 2022. [1](#), [2](#), [3](#), [4](#), [5](#)
- [61] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. [3](#)
- [62] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. [2](#)
- [63] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Wei Wei, Tingbo Hou, Yael Pritch, Neal Wadhwa, Michael Rubinstein, and Kfir Aberman. HyperDreamBooth: HyperNetworks for Fast Personalization of Text-to-Image Models, July 2023. arXiv:2307.06949 [cs]. [2](#), [3](#)
- [64] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative Radiance Fields for 3D-Aware Image Synthesis. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20154–20166. Curran Associates, Inc., 2020. [1](#)
- [65] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9243–9252, 2020. [7](#)
- [66] James Seale Smith, Yen-Chang Hsu, Lingyu Zhang, Ting Hua, Zsolt Kira, Yilin Shen, and Hongxia Jin. Continual Diffusion: Continual Customization of Text-to-Image Diffusion with C-LoRA, Apr. 2023. arXiv:2304.06027 [cs]. [3](#)
- [67] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gerard Medioni. Regressing Robust and Discriminative 3D Morphable Models with a Very Deep Neural Network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1493–1502, Honolulu, HI, July 2017. IEEE. [2](#)
- [68] Luan Tran and Xiaoming Liu. On learning 3d face morphable model from in-the-wild images. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):157–171, 2019. [2](#)
- [69] Alex Trevithick, Matthew Chan, Michael Stengel, Eric Chan, Chao Liu, Zhiding Yu, Sameh Khamis, Manmohan Chandraker, Ravi Ramamoorthi, and Koki Nagano. Real-time radiance fields for single-image portrait view synthesis. *ACM Transactions on Graphics (TOG)*, 42(4):1–15, 2023. [1](#), [4](#), [5](#)
- [70] Neal Wadhwa, Rahul Garg, David E Jacobs, Bryan E Feldman, Nori Kanazawa, Robert Carroll, Yair Movshovitz-Attias, Jonathan T Barron, Yael Pritch, and Marc Levoy. Synthetic depth-of-field with a single-camera mobile phone. *ACM Transactions on Graphics (ToG)*, 37(4):1–13, 2018. [6](#)
- [71] Hai Wang, Xiaoyu Xiang, Yuchen Fan, and Jing-Hao Xue. Customizing 360-Degree Panoramas through Text-to-Image Diffusion Models, Nov. 2023. arXiv:2310.18840 [cs]. [3](#)
- [72] Mika Westerlund. The emergence of deepfake technology: A review. *Technology innovation management review*, 9(11), 2019. [2](#)
- [73] Jiaxin Xie, Hao Ouyang, Jingtian Piao, Chenyang Lei, and Qifeng Chen. High-fidelity 3d gan inversion by pseudo-multi-view optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 321–331, 2023. [1](#)

- [74] Ziyang Xie, Junge Zhang, Wenye Li, Feihu Zhang, and Li Zhang. S-NeRF: Neural Radiance Fields for Street Views, Mar. 2023. arXiv:2303.00749 [cs]. 5
- [75] Shunyu Yao, RuiZhe Zhong, Yichao Yan, Guangtao Zhai, and Xiaokang Yang. DFA-NeRF: Personalized Talking Head Generation via Disentangled Face Attributes Neural Rendering, Jan. 2022. arXiv:2201.00791 [cs]. 2
- [76] Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A. Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, Sergey Yekhanin, and Huishuai Zhang. Differentially Private Fine-tuning of Language Models, July 2022. arXiv:2110.06500 [cs, stat]. 2
- [77] Ziyang Yuan, Yiming Zhu, Yu Li, Hongyu Liu, and Chun Yuan. Make Encoder Great Again in 3D GAN Inversion through Geometry and Occlusion-Aware Encoding, Mar. 2023. arXiv:2303.12326 [cs]. 8
- [78] Libing Zeng, Lele Chen, Yi Xu, and Nima Khademi Kalantari. Mystyle++: A controllable personalized generative prior. In *SIGGRAPH Asia 2023 Conference Papers, SA '23*, New York, NY, USA, 2023. Association for Computing Machinery. 3, 6
- [79] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3836–3847, October 2023. 3
- [80] Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adaptive Budget Allocation for Parameter-Efficient Fine-Tuning, Mar. 2023. arXiv:2303.10512 [cs]. 3, 8
- [81] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, Salt Lake City, UT, June 2018. IEEE. 4
- [82] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I. Chang, and Yan Xu. Large Scale Image Completion via Co-Modulated Generative Adversarial Networks, Mar. 2021. arXiv:2103.10428 [cs]. 6