

Controllable 3D Face Generation with Conditional Style Code Diffusion

Xiaolong Shen^{1,2*}, Jianxin Ma², Chang Zhou², Zongxin Yang^{1†}

¹ ReLER, CCAI, Zhejiang University, China

² Alibaba Group, China

{sxlongcs, zongxinyang}@zju.edu.cn, {majx13fromthu,ericzhou.zc}@alibaba-inc.com

Abstract

Generating photorealistic 3D faces from given conditions is a challenging task. Existing methods often rely on time-consuming one-by-one optimization approaches, which are not efficient for modeling the same distribution content, *e.g.*, faces. Additionally, an ideal controllable 3D face generation model should consider both facial attributes and expressions. Thus we propose a novel approach called TEx-Face (TExT & Expression-to-Face) that addresses these challenges by dividing the task into three components, *i.e.*, 3D GAN Inversion, Conditional Style Code Diffusion, and 3D Face Decoding. For 3D GAN inversion, we introduce two methods which aim to enhance the representation of style codes and alleviate 3D inconsistencies. Furthermore, we design a style code denoiser to incorporate multiple conditions into the style code and propose a data augmentation strategy to address the issue of insufficient paired visual-language data. Extensive experiments conducted on FFHQ, CelebA-HQ, and CelebA-Dialog demonstrate the promising performance of our TEx-Face in achieving the efficient and controllable generation of photorealistic 3D faces. The code will be available at <https://github.com/sxl142/TEx-Face>.

Introduction

Face generation, as a longstanding task in image synthesis, is extensively utilized in artistic creation and portrait generation. Benefits from well-designed 2D/3D unconditional generation frameworks (Karras et al. 2020; Chan et al. 2022), current face generation models have achieved remarkable photorealism in both 2D and 3D. However, introducing multiple conditional representations (Yang, Zhuang, and Pan 2021) into face generation, particularly in the context of 3D, remains largely unexplored.

Existing Text-to-3D generation techniques (Poole et al. 2022; Huang et al. 2023a; Chen et al. 2023; Xu, Yang, and Yang 2023) typically involve training an initialized NeRF (Mildenhall et al. 2020) model to generate a specific object using pre-trained text-to-image diffusion models (Romach et al. 2022). Although these approaches offer flexibility and can produce various objects, **generating a large**

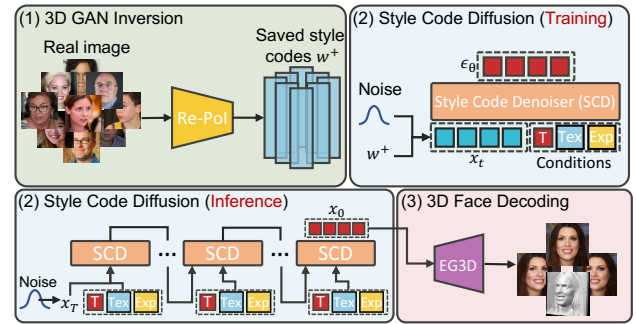


Figure 1: An overview of our pipeline. We train an inversion model called Re-PoI and save the style codes it infers. These saved codes are then used to train a style code denoiser with three conditions, *i.e.*, time steps, text prompts, and expression codes. When inference, we decode the generated style codes into 3D faces using EG3D.

number of objects from the same distribution is time-consuming, such as generating 3D faces with one-by-one optimization schema. Besides, some text-to-3D face generation techniques (Zhang et al. 2023; Wu et al. 2023) primarily focus on generating texture for parametric face models, *e.g.*, 3D Morphable face model (Banz and Vetter 1999). While parametric models make it possible to explicitly control the expression, pose, and lighting of faces, the generated results often **lack realism** since the parametric geometry lacks finer details such as eyes, teeth, and hair.

Human faces consist of two main components, *i.e.*, appearance and skeleton. The appearance represents the visual characteristics that differentiate individuals, while the skeleton is responsible for different poses and expressions. From this perspective, an ideal controllable generative model should be able to understand these inherent facial properties. Thus, we need to find a reasonable way to represent these properties for the generative model. The emergence of CLIP (Radford et al. 2021) has bridged the gap between images and text, enabling successful text-conditioned generation like (Nichol et al. 2022; Ramesh et al. 2022; Saharia et al. 2022). This motivates us to choose the text to represent facial appearances. Regarding the other component skeleton, a typical representation is a set of keypoints. Alternatively, we can

*Xiaolong Shen worked on this at his Alibaba internship.

†Zongxin Yang is the corresponding author.

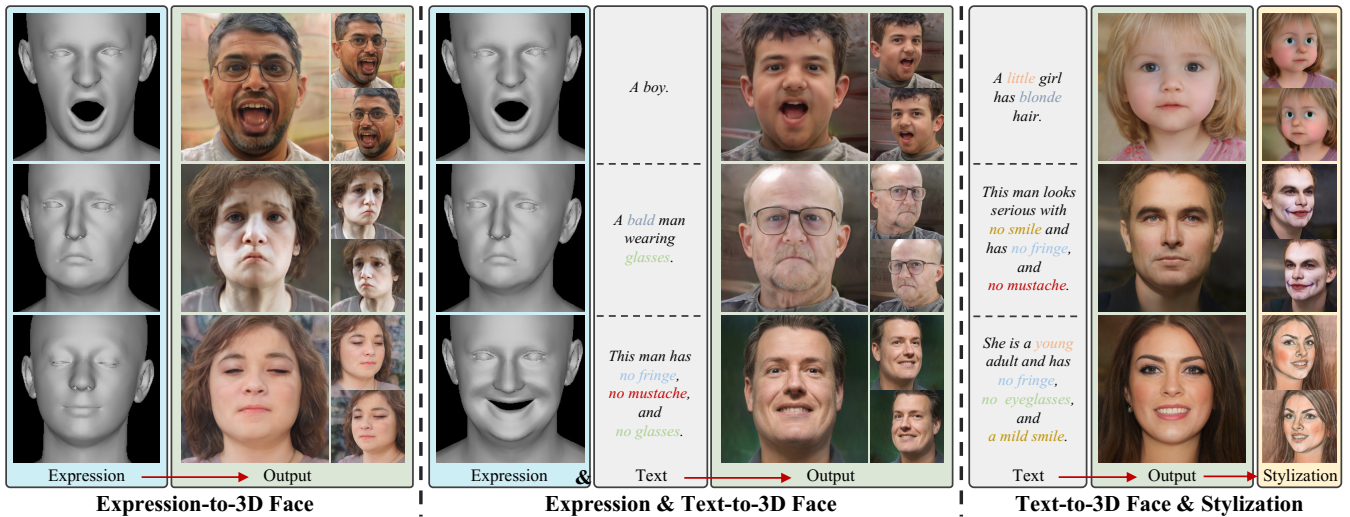


Figure 2: Our pipeline enables conditional 3D face generation using text, expression codes, and both of them. The stylization is achieved by StyleGAN-NADA (Gal et al. 2021), a GAN-based style transfer method.

leverage the parameters of the 3D Morphable face model, *i.e.*, expression code. Compared to keypoints, these parameters can yield 3D meshes through the parametric model, offering greater flexibility in controlling the generated faces.

Based on the aforementioned observations, we aim to explore an approach to address two key aspects: **1) generating photorealistic 3D faces efficiently**, and **2) controllable generation based on the given conditions**, *i.e.*, text prompts and expression codes. First, we consider that the current unconditional 3D GAN model (Chan et al. 2022) is a suitable foundation model, as it has demonstrated the ability to produce convincing 3D faces efficiently. Subsequently, choosing this model requires considering how to equip the 3D GAN with text-conditioned and expression-conditioned abilities. Upon analyzing the 3D GAN framework, we observe that the produced results are primarily influenced by style codes. Therefore, we resort to GAN Inversion technology (Tov et al. 2021; Alaluf, Patashnik, and Cohen-Or 2021) to obtain the necessary style codes. Since there is a lack of encoder-based 3D GAN Inversion, we need to extend 2D inversion methods to the 3D domain. However, simply extending the 2D methods can lead to overfitting to the given camera pose, yielding undesirable novel view rendering, as shown in Figure 3. Second, we need to consider how to inject multi-conditions into the obtained style codes. With the recent surge in diffusion-based generative models (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2022; Zhou, Yang, and Yang 2023) applied to images, there is great interest in exploring the potential of diffusion models for learning style codes, given their ability to facilitate flexible condition injections. However, aligning the given conditions with semantic-agnostic style codes remains a challenging open problem.

To address these issues, we propose a novel framework named TEx-Face, as shown in Figure 1, involving three components. **1) 3D GAN Inversion.** We design two methods for 3D GAN Inversion, *i.e.*, Pose-guided Inversion (PoI) Pretrain-

ing and Refined Pose-guided Inversion (Re-PoI) Finetuning, targeting to alleviate the 3D inconsistency and enhance the representation of style codes. In pretraining, we expect that the images’ style codes are consistent under the different camera views for one identity. Thus, we need to project each view’s style codes onto a canonical style code. Here, we choose style codes under the front view as canonical codes due to the largest visible facial appearance. Specifically, due to the lack of multi-view images, we first generate synthetic canonical style codes using the Mapping Network in 3D GAN and then use these codes to render images under randomly sampled camera poses. After that, we can leverage the PoI to learn a mapping that projects the different views’ style codes onto one canonical style code, yielding a view-invariant code. In the finetuning stage, we apply coarse-to-fine schema, incorporating a new branch to learn the difference between the reconstructed and input images and complement details for high-quality reconstruction. **2) Style Code Diffusion.** We introduce a style codes denoiser to decoupled inject the text embedding extracted by CLIP (Radford et al. 2021) and expression code extracted by EMOCA (Danecek, Black, and Bolkart 2022) into the style codes for learning the distribution alignment. Moreover, considering the insufficient of paired visual-language face data, we propose a data augmentation strategy for improving the diversity of the origin dataset, which utilizes a manipulation method and an image captioning method to automatically generate paired data. **3) 3D Face Decoding.** The style codes are crucial for this task, they directly affect the quality of the generated results. Hence, we select the most advanced method, EG3D (Chan et al. 2022), as our decoding model. This method is based on the architecture of StyleGAN2 (Karras et al. 2020), which results in a well-behaved style code space.

We conduct extensive experiments on several datasets, *i.e.*, FFHQ (Abdal, Qin, and Wonka 2019), CelebA-HQ (Karras et al. 2018), and CelebA-Dialog (Jiang et al. 2021). Experi-




Figure 3: Simply extending the 2D Inversion method (Hu et al. 2022) leads to bad novel view synthesis, especially when inputting an image of side views.

ments show that our method improves the image quality and alleviates the 3D inconsistency for 3D GAN Inversion, as shown in Figure 3. To evaluate the 3D Face generation quality, we provide various metrics, including FID, CLIP score, expression score, and user studies. The results indicate that our model is capable of producing photorealistic 3D faces that align with the given text prompts and expression codes.

Overall, our contributions are summarized as follows,

- To our best knowledge, we make the first attempt to explore controllable 3D face generation using multi-conditions and one-time optimization method, which decomposes it into three parts, *i.e.*, 3D GAN Inversion, Style Code Diffusion, and 3D Face Decoding.
- We design two methods for 3D GAN Inversion, *i.e.*, Pose-guided Inversion Pretraining and Refined Pose-guided Inversion Funetuning, to enhance the style codes and alleviate the 3D appearance inconsistency.
- We introduce a style code denoiser to inject text embedding and expression code into the style codes. Moreover, we propose a data augmentation strategy to generate paired data automatically.
- Extensive experiments show our framework achieves promising conditional 3D face generation results.

Related Work

2D&3D GAN Inversion

GAN Inversion (Zhu et al. 2016; Tov et al. 2021; Roich et al. 2022; Wang et al. 2022a; Hu et al. 2022; Xie et al. 2022; Yin et al. 2023) focuses on finding a semantically rich latent space for the images, which can be further used to edit the images via some latent space manipulation methods (Wei et al. 2022; Patashnik et al. 2021; Zhu et al. 2022; Shen et al. 2020; Zhu et al. 2020). 2D GAN Inversion method are mainly based on StyleGAN2 (Karras et al. 2020) owing to the well-designed architecture. There are three types of Inversion methods, *i.e.*, optimize-based (Xie et al. 2022; Yin et al. 2023; Wu, Lischinski, and Shechtman 2021), encoder-based (Tov et al. 2021; Wang et al. 2022a; Hu et al. 2022; Alaluf, Patashnik, and Cohen-Or 2021), and hybrid-based methods (Alaluf, Patashnik, and Cohen-Or 2021; Chai et al.

2021; Ko et al. 2023). Optimize-based methods directly optimize a latent code for a specific image, which makes the reconstructed images more realistic and fidelity but usually takes a long time. Encoder-based methods only spend time on training, while the quality of reconstructed images is worse than optimize-based methods. Hybrid-based methods leverage the latent codes inferred by encoder-based methods as initial latent codes for optimize-based methods. However, these method is still time-consuming. Recently, EG3D (Chan et al. 2022) leverages StyleGAN2 generator as a backbone to yield tri-plane for 3D face generation, which makes it possible that can simply extend 2D GAN Inversion methods by changing the StyleGAN2 with EG3D due to inherent similarity. However, simply extending 2D methods results in 3D inconsistency. While some works (Yin et al. 2023; Xie et al. 2022; Ko et al. 2022) target to solve this problem, these methods are optimization-based, which is not suitable for generating a large number of style codes.

Conditional Generation

2D generation. Earlier works (Xu et al. 2018; Yang et al. 2019) are mainly constructed on Generative Adversarial Networks (GANs) (Goodfellow et al. 2014). The recent surge in Diffusion models (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2022; Zhou, Yang, and Yang 2023) is a class of generative models that excel at generating realistic images through iterative refinement of a sampled noise. DDPM (Ho, Jain, and Abbeel 2020) employs a dual Markov chain approach for unconditional image generation. Imagen (Saharia et al. 2022) explores large language models (Raffel et al. 2020) trained on text data as text encoders for text-to-image diffusion models. Unlike above, stable-diffusion (Rombach et al. 2022) applies the diffusion paradigm to the latent space of a Variational autoencoder (VAE) and introduces cross-attention layers into UNet for additional conditions.

3D generation. In recent years, the field of text-to-image generative modeling (Nichol et al. 2022; Wang et al. 2022b; Rombach et al. 2022) has achieved significant advancements, leading to a surge of interest in text-to-3D generation (Poole et al. 2022; Huang et al. 2023a; Xu, Yang, and Yang 2023). CLIP-Forge (Sanghi et al. 2022) capitalizes on the shared text-image embedding space of CLIP to train on image embeddings and utilize text embeddings for inference, thereby enabling text-to-3D capabilities. DreamFusion (Poole et al. 2022) employs a pretrained text-to-image diffusion model as a potent image prior and introduces a Score-Distillation loss to supervise the generation of 3D objects represented by Neural Radiance Fields (NeRFs (Mildenhall et al. 2020)). Following DreamFusion, AvartarFusion (Huang et al. 2023a) and SEEAvatar (Xu, Yang, and Yang 2023) focus on high-quality 3D human generation by combining parametric human models and text-to-image diffusion models. For text-to-3D face generation, current methods (Zhang et al. 2023; Wu et al. 2023) mainly build on a 3D Morphable face model and generate 3D faces with geometry and texture. Owing to the parametric model, these models can explicitly control the expression, pose, and so on. However, the generation results usually lack details, *e.g.*, hair, and look unnatural.

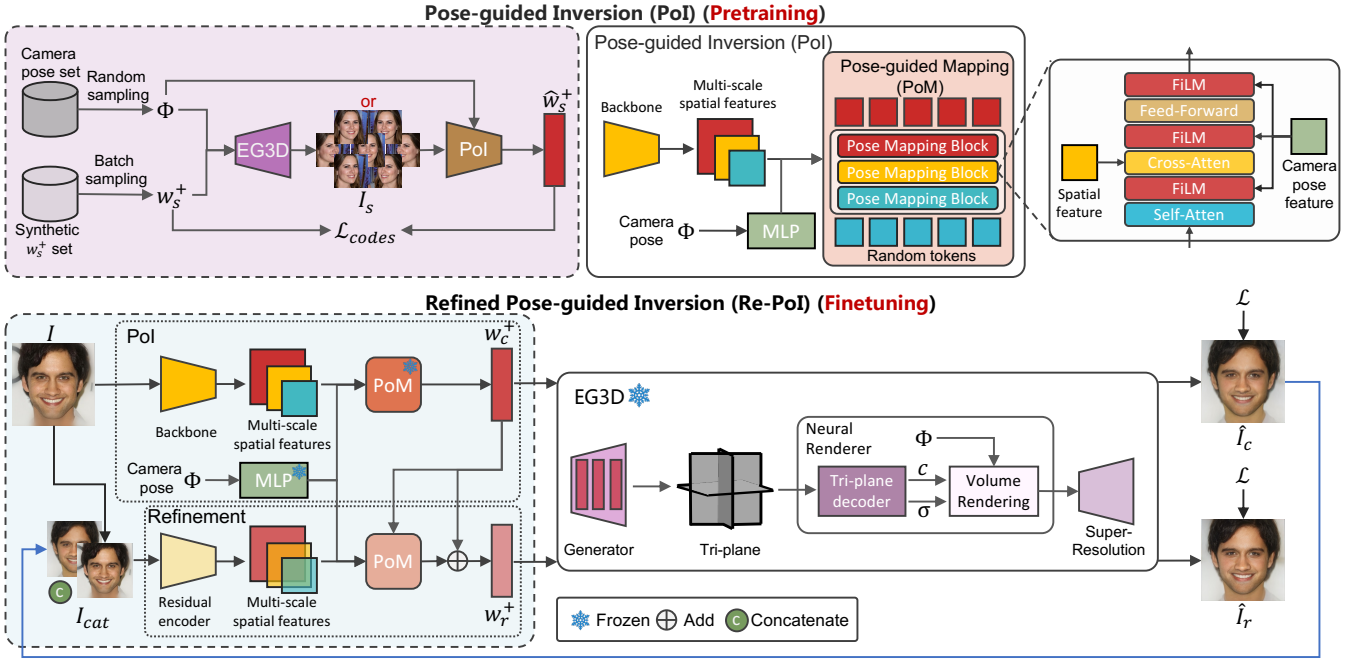


Figure 4: An overview of our inversion method. In pretraining, we first leverage the synthetic multi-view images to learn a mapping that projects the style codes under different views onto one style code, which yields a view-invariant style code, thereby alleviating 3D inconsistency. In finetuning, we freeze the learned mapping (PoM) and MLP in the PoI and further append the refinement branch to improve the quality of the style codes. Note the PoM of the refinement branch is used for training.

Method

In this section, we first recap the EG3D (Chan et al. 2022) and diffusion models in preliminary, and then we introduce the designed inversion model and style code diffusion model.

Preliminary

EG3D. As shown in Figure 4, EG3D is composed of a Mapping Network and a Synthesis Network. The Mapping Network generates a style code w^+ , $w^+ \in \mathbf{R}^{14 \times 512}$ by feeding a combination of a sampled Gaussian vector z , $z \in \mathbf{R}^{512}$ and a camera pose ϕ , $\phi \in \mathbf{R}^{25}$ sampled from datasets. Here we only draw the Synthesis Network because the inversion process replaces the Mapping Network to yield style codes. After that, the Synthesis Network produces a Tri-plane representation via a series of convolution layers modulated by the style code w^+ . The resulting Tri-plane representation is then decoded by the Neural Renderer, producing coarse images. Finally, the generated coarse images are refined by a Super-Resolution Network, enhancing their quality.

Diffusion models. Recently, diffusion models (Ho, Jain, and Abbeel 2020; Ho and Salimans 2022; Song, Meng, and Ermon 2022) have demonstrated remarkable capabilities in conditional or unconditional generation tasks. Diffusion models are a class of probabilistic models that involve two stages, *i.e.* forward and reverse processes. The forward process works by gradually adding noise to the data in T steps, generating a sequence of noise samples x_1, \dots, x_T . It follows the rule of the Markov chain. The resulting noise sample x_t can be

defined as follows,

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}) \quad (1)$$

$$x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (2)$$

where $\alpha_t = \prod_{i=1}^t (1 - \beta_i)$, $\{\beta_t\}_{t=1}^T$ are fixed or learned variance schedule. In the reverse stage, we need to learn a model $p_\theta(x_{t-1}|x_t)$ to approximate the conditional probabilities $q(x_{t-1}|x_t)$ because we cannot easily estimate the distribution of the whole dataset. Ho *et al.* find the training objective can be simplified to make the diffusion model learn better. The function is as follows,

$$\min_{\theta} \mathbb{E}_{x_0 \sim q(x_0), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), t \sim [1, T]} \|\epsilon - \epsilon_\theta(x_t, t)\|_2^2. \quad (3)$$

3D GAN Inversion

Pose-guided Inversion (PoI) Pretraining. We demonstrate that simply utilizing the 2D inversion method can result in overfitting to the input view as shown in Figure 3, leading to 3D inconsistency when rendering images from novel views. To address this, we propose an explicit learning method that aims to project style codes under different views onto a single style code, producing a view-invariant style code. We consider a style code under the front view as canonical code. However, we lack multi-view images of one identity to model this mapping. To achieve this, we leverage the pretrained EG3D to produce synthetic multi-view images by using this canonical style code as depicted in Figure 4.

Specifically, we first sample 50K style codes w_s^+ from the Mapping Network in EG3D and save them. Subsequently,

we batch sample w_s^+ which are then used to render images with randomly sampled camera poses Φ . Lastly, the produced images will be encoded as style codes \hat{w}_s^+ through our proposed PoI. The PoI consists of a backbone and a Pose-guided Mapping Network (PoM). Given a synthetic image I_s , the backbone extracts multi-scale features M from this image, and then the features with the camera feature will be fed into the PoM. The PoM has two important layers, *i.e.*, a cross-attention layer and a Feature-wise Linear Modulation (FiLM) layer. The cross-attention layer is used for decoding multi-scale features, while the FiLM is responsible for mapping these decoded features to the expected canonical space by conditioning the camera feature. The main processes are as follows,

$$\begin{aligned} \text{CrossAtten}(Q^c, K_{M_i^c}, V_{M_i^c}) &= \text{Softmax}\left(\frac{Q^c K_{M_i^c}^T}{\sqrt{C}}\right) V_{M_i^c}, \\ \text{FiLM}(f_p, f_c) &= \gamma(f_c) f_p + \beta(f_c), \end{aligned} \quad (4)$$

where M_i^c is an i_{th} -scale spatial feature, γ and β are Linear function. The camera feature f_c is extracted by the Multilayer Perceptron (MLP). The f_p is the feature of the previous layer. Finally, The PoM transforms the random tokens to the coarse style codes w_c^+ by several Pose Mapping Blocks.

Refined Pose-guided Inversion (PoI) Finetuning. In pre-training, we train the PoI on synthetic data. Thus we need to finetune the PoI on real data. However, fully finetuning the model will discard the learned projection ability that maps the style codes of multi-view images to the canonical style code. Therefore, we freeze the PoM and only finetune the backbone. Besides, simply finetuning merely yields coarse codes, lacking some details. Hence, we propose to complement details by using a coarse-to-fine schema. Similar to the PoI, the refined module involves a PoM and residual encoder.

Specifically, we concatenate the coarse image \hat{I}_c and ground truth \hat{I} as the input I_{cat} of the Residual Encoder (RE). Similarly, the encoder extracts multi-scale features M_r for the Re-PoM. In the first block of the Re-PoM, it will use the coarse style codes as input to query the extracted residual multi-scale features. Finally, the output of Re-PoM will be added with the coarse style codes as the final results w_r^+ . The main processes are formulated as follows.

$$M_r = RE(\text{concat}(I, \hat{I}_c)),$$

$$\text{CrossAtten}(Q_r, K_{M_r^c}, V_{M_r^c}) = \text{Softmax}\left(\frac{Q_r K_{M_r^c}^T}{\sqrt{C}}\right) V_{M_r^c},$$

$$w_r^+ = w_c^+ + \text{RePoM}(M_r, w_c^+). \quad (5)$$

Training objective. To learn a view-invariant style code in the pretraining stage, we apply loss on the style codes.

$$\mathcal{L}_{codes} = \|w_s^+ - \hat{w}_s^+\|_2 \quad (6)$$

In finetuning, we apply MSE, LPIPS, and ID loss for reconstructed images. Besides, inspired by (Abdal et al. 2023), we apply the depth loss on the reconstructed images, which will further improve the quality of geometry.

$$\mathcal{L} = \lambda_{rec} * \mathcal{L}_{rec} + \lambda_{lips} * \mathcal{L}_{lips} + \lambda_{id} * \mathcal{L}_{id} + \lambda_{dep} * \mathcal{L}_{dep} \quad (7)$$

Conditional Style Code Diffusion

Considering the facial ingredients, *i.e.*, appearance and expression, we expect the model can be controlled by these two

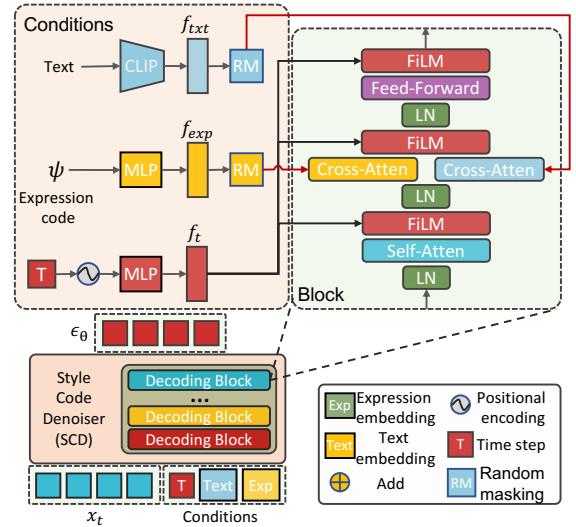


Figure 5: A schematic of Style Code Denoiser.

aspects. Thus we propose incorporating text embedding and expression codes to guide the generation of appearance and skeleton, respectively. However, incorporating these conditions into the model flexibly and effectively poses a challenge. Upon analyzing the overall structure of the EG3D, we identify the style code inferred by the Mapping Network of the EG3D as a crucial component in determining the generated results. Thus we design a style code denoiser to inject these conditions into this style code.

Style Code Denoiser. To achieve a multi-conditioned model, we structurally decouple the conditions by employing a separated cross-attention layer. Concretely, as shown in Figure 5, there are three conditions, *i.e.*, text, expression code, and time step. To model the text embedding f_{txt} , we utilize the CLIP text encoder. CLIP enables us to encode textual information into a meaningful representation that can guide the generation of facial appearance. For the expression embedding, we first employ EMOCA (Danecek, Black, and Bolkart 2022) to predict a set of expression code, denoted as ψ , $\psi \in \mathbf{R}^{50}$ and then utilize an MLP to obtain the final representation f_{exp} . To incorporate the time step information, we utilize positional encoding and an MLP, which allows us to represent the temporal aspect of the generation process. Next, we apply random masking to both text and expression embeddings (f_{txt} , f_{exp}) to implement classifier-free guidance. Then these features are used as keys and values for extracting information through a separated cross-attention layer in the decode block. The time embedding is responsible for transforming the results to the given time steps by using FiLM layers. Since we need to inject two conditions into the model, the training objective will be changed with multi-conditioned as follows.

$$\min_{\theta} \mathbb{E}_{x_0 \sim q(x_0), \epsilon \sim \mathcal{N}(0, \mathbf{I}), t \sim [1, T]} \|\epsilon - \epsilon_{\theta}(x_t, t, c_{txt}, c_{exp})\|_2^2. \quad (8)$$

Data augmentation. Existing visual-language face dataset Celeba-Dialog (Jiang et al. 2021) is built on the Celeba-HQ (Karras et al. 2018), which lacks some attribute annotation, *e.g.*, hair color or hairstyle, and mainly consists of adults, lacking children and old people. Based on these observations,

Method	MSE ↓	LPIPS ↓	ID ↓	FID ↓
e4e	0.0354	0.1970	0.1903	33.97
ST	0.0259	0.1677	0.2199	34.59
Ours	0.0252	0.1721	0.1788	30.15

Table 1: GAN Inversion.

Method	Type	FID ↓	CLIP score ↑
TediGAN	2D	157.81	24.27
CD	2D	111.36	24.51
Ours	3D	52.91	24.83

Table 3: Conditional face generation.

	MSE ↓		LPIPS ↓		ID ↓	
	Self	Others	Self	Others	Self	Others
ST	0.014	0.168	0.107	0.477	0.175	0.375
Ours	0.012	0.048	0.105	0.189	0.163	0.244

Table 2: Evaluation of 3D consistency.

	Text	Expression
Acc(%) ↑	79.3	60.0

Table 4: User studies.

	Happy	Sad	Surprise	Disgust	Anger
Confidence ↑	91.49	53.97	90.65	35.92	34.26

Table 5: Expression score of our method.

we try to bridge these gaps by using existing manipulation methods and image captioning methods to automatically generate paired data. Specifically, we implement HairCLIP (Wei et al. 2022) on the CelebA-HQ training set to obtain 2D text-paired images of different hair colors and hairstyles. Then we leverage our method to inverse these images to obtain style codes. Moreover, we use captioning model (Li et al. 2023) to generate text for FFHQ (Abdal, Qin, and Wonka 2019). The benefits of using captioning models are that we can use more face data, thus the trained model can handle a wider face distribution and improve facial diversity. For text-guided manipulation methods, we can use them to manipulate face images to obtain more facial attributes paired with text.

Implementation Details

We use Adam (Kingma and Ba 2017) optimizer with linear warm-up and Cosine Annealing (Loshchilov and Hutter 2017) scheduler. We set the loss weights as follows: $\lambda_{rec} = 1$, $\lambda_{lips} = 0.8$, $\lambda_{id} = 0.2$. We use four Nvidia Tesla V100 (16G) with batch size 8 to train our inversion model, and with batch size 256 for style code diffusion. For data processing, we follow EG3D, using off-the-shelf pose estimators to extract approximate camera extrinsic.

Experimental Results

Datasets & Metrics. We train our inversion model on FFHQ (Abdal, Qin, and Wonka 2019) and test it on CelebA-HQ (Karras et al. 2018) test set. We use CelebA-Dialog (Jiang et al. 2021) and some data processed by our proposed data augmentation strategy to train our diffusion model. We evaluate inversion methods by Mean Squared Error (MSE), Perceptual Similarity (LPIPS (Zhang et al. 2018)), FID (Heusel et al. 2017), and identity similarity (ID). For generation quality, we use FID, CLIP score, expression score, and user studies.

Setting. Optimization-based GAN inversion methods (Roich et al. 2022; Xie et al. 2022) are time-consuming when applied to obtain a large number of style codes. Thus we only focus on encoder-based inversion methods. Due to the lack of released encoder-based 3D inversion methods, we reimplement 2D state-of-the-art inversion methods (Tov et al. 2021; Hu et al. 2022) for 3D GAN inversion by simply replacing StyleGAN2 with EG3D.

Quantitative Comparison.

GAN Inversion. As shown in Table 1, Style Transformer (ST) (Hu et al. 2022) obtains the lowest LPIPS metric but other metrics are worse than our method. Moreover, ST suffers from 3D inconsistency, as shown in Figure 3. To evaluate the 3D consistency, we conducted experiments on the MEAD (Wang et al. 2020) dataset in Table 2. “Self” measures self-reconstruction. “Others” measures novel view synthesis, which means rendering images under different views with the input. MEAD has different view images of one identity. Thus, we can compute metrics with the ground truth. It shows our model outperforms the ST, **particularly in “others”**. It also demonstrates the 2D method is easier to overfit to the input view, leading to worse performance in novel view synthesis.

Conditional 3D Face Generation. Due to a lack of similar methods, we compare our method with 2D conditional generation methods (Xia et al. 2021; Huang et al. 2023b) using the FID metric for generation quality and CLIP score for text-image matching. The results, presented in Table 3, demonstrate that our method achieves the best performance. Furthermore, we evaluate our expression-to-3D faces by choosing five emotions (Happy, Sad, Surprise, Disgust, and Anger) and use the recognition model of EMOCA (Danecek, Black, and Bolkart 2022) to test the generation results. The average confidence of the corresponding emotion class is provided in Table 5. Our model achieves high confidence for Happy and Surprise expressions. While Sad, Disgust, and Anger do not reach high confidence, they are significantly better than a random probability $1/11 * 100 = 9\%$ since the recognition model has 11 classes. Moreover, we consider the data imbalance also influences the performance. As shown in Table 4, user studies are provided to validate text-image matching and expression-image matching. The results show that our model can generate plausible 3D faces based on given conditions.

Inference time. PTI (Roich et al. 2022), an optimization-based GAN inversion method, requires about 120s to process a single sample. Consequently, employing this method for a large number of samples is time-consuming. In contrast, our



Figure 6: Visual comparison of Text-to-Face.

Type	Expression score \uparrow	CLIP score \uparrow
Coupled	58.38	20.62
Decoupled	59.70	21.86

Table 6: Different types of cross-attention layers.



Figure 7: Visual comparison with Coupled and Decoupled cross-attention layers.

method takes only 0.10 seconds per sample. Regarding style code diffusion, we employ a DDIM sampler with a sampling step of 100. This process takes approximately 5.1s.

Qualitative Comparison.

3D GAN inversion. In Figure 3, Our method shows improvements in facial symmetry and overall geometric accuracy compared with simply extending the 2D method to a 3D domain. More results are placed in the Supplementary.

Conditional 3D face generation. As illustrated in Figure 2, we provide the results of text-to-3D face, expression-to-3D face, and text & expression-to-3D face, which demonstrate that our pipeline is capable of generating plausible 3D faces. Moreover, we utilize StyleGAN-NADA (Gal et al. 2021), a GAN-based style transfer method, which enables us to generate 3D faces in different domains by simply replacing the original EG3D model with a style-EG3D model trained using this method. For visual comparison with other methods, we show it in Figure 6, our pipeline surpasses other methods both in quality and text-face matching.

Ablation Study.

3D GAN inversion. We provide the results without refinement in the Re-POI finetuning stage in Table 7. It demon-

Method	MSE \downarrow	LPIPS \downarrow	ID \downarrow	FID \downarrow
w/o refinement	0.0404	0.2128	0.2156	35.90
w/ refinement	0.0252	0.1721	0.1788	30.15

Table 7: Ablation studies of Re-PoI.

strates that the refinement module significantly improves the performance. More results are in the supplementary.

Conditional 3D face generation. We conduct experiments that apply different cross-attention layers as shown in Table 6 and Figure 7. Coupled means using the same cross-attention to model text and expression embedding, while Decoupled represents using separated cross-attention. The results show that the decoupled layer yields better performance and produces 3D faces that align well with the given conditions.

Conclusion

We propose TEx-Face for conditional 3D face generation, consisting of three components, *i.e.*, 3D GAN inversion, Conditional Style Code Diffusion, and 3D Face Decoding. Under the framework, we design two methods for 3D GAN inversion to enhance the representation of style codes and alleviate the 3D inconsistency. For the multi-conditioned generation, we propose a style code denoiser that successfully decoupled models three conditions, *i.e.*, time step, text, and expression code. Moreover, we introduce a data augmentation strategy to generate paired data automatically. Based on these designs, our pipeline shows remarkable results.

Limitations and Future Work. The generation quality is highly dependent on the capability of EG3D and the inversion model. Moreover, the training data are imbalanced in expression. Thus, our model may focus on the majority of expressions, *e.g.*, happy. The data augmentation strategy relies on manipulation and captioning methods. Hence, the ability of these models dominates the quality of the generated paired data. We will extend this framework to some methods (An et al. 2023; Gao et al. 2022) having a similar architecture with EG3D and will explore data augmentation strategy with our method in face detection (*e.g.*, (Zhou et al. 2021)).

Acknowledgements

This work was supported by the Fundamental Research Funds for the Central Universities (No. 226-2023-00051).

References

- Abdal, R.; Lee, H.-Y.; Zhu, P.; Chai, M.; Siarohin, A.; Wonka, P.; and Tulyakov, S. 2023. 3davatar: Bridging domains for personalized editable avatars. In *CVPR*, 4552–4562.
- Abdal, R.; Qin, Y.; and Wonka, P. 2019. Image2stylegan: How to embed images into the stylegan latent space? In *ICCV*, 4432–4441.
- Alaluf, Y.; Patashnik, O.; and Cohen-Or, D. 2021. Restyle: A residual-based stylegan encoder via iterative refinement. In *ICCV*, 6711–6720.
- An, S.; Xu, H.; Shi, Y.; Song, G.; Ogras, U. Y.; and Luo, L. 2023. PanoHead: Geometry-Aware 3D Full-Head Synthesis in 360deg. In *CVPR*, 20950–20959.
- Blanz, V.; and Vetter, T. 1999. A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 187–194.
- Chai, L.; Zhu, J.-Y.; Shechtman, E.; Isola, P.; and Zhang, R. 2021. Ensembling with deep generative views. In *CVPR*, 14997–15007.
- Chan, E. R.; Lin, C. Z.; Chan, M. A.; Nagano, K.; Pan, B.; De Mello, S.; Gallo, O.; Guibas, L. J.; Tremblay, J.; Khamis, S.; et al. 2022. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 16123–16133.
- Chen, R.; Chen, Y.; Jiao, N.; and Jia, K. 2023. Fantasia3D: Disentangling Geometry and Appearance for High-quality Text-to-3D Content Creation. arXiv:2303.13873.
- Danecek, R.; Black, M. J.; and Bolkart, T. 2022. EMOCA: Emotion Driven Monocular Face Capture and Animation. In *CVPR*, 20311–20322.
- Gal, R.; Patashnik, O.; Maron, H.; Chechik, G.; and Cohen-Or, D. 2021. StyleGAN-NADA: CLIP-Guided Domain Adaptation of Image Generators. arXiv:2108.00946.
- Gao, J.; Shen, T.; Wang, Z.; Chen, W.; Yin, K.; Li, D.; Litany, O.; Gojcic, Z.; and Fidler, S. 2022. GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images. arXiv:2209.11163.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In *NIPS*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NIPS*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *NIPS*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-Free Diffusion Guidance. arXiv:2207.12598.
- Hu, X.; Huang, Q.; Shi, Z.; Li, S.; Gao, C.; Sun, L.; and Li, Q. 2022. Style transformer for image inversion and editing. In *CVPR*, 11337–11346.
- Huang, S.; Yang, Z.; Li, L.; Yang, Y.; and Jia, J. 2023a. AvatarFusion: Zero-shot Generation of Clothing-Decoupled 3D Avatars Using 2D Diffusion. In *ACM MM*, 5734–5745.
- Huang, Z.; Chan, K. C. K.; Jiang, Y.; and Liu, Z. 2023b. Collaborative Diffusion for Multi-Modal Face Generation and Editing. arXiv:2304.10530.
- Jiang, Y.; Huang, Z.; Pan, X.; Loy, C. C.; and Liu, Z. 2021. Talk-to-edit: Fine-grained facial editing via dialog. In *ICCV*, 13799–13808.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *ICLR*.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *CVPR*.
- Kingma, D. P.; and Ba, J. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980.
- Ko, J.; Cho, K.; Choi, D.; Ryoo, K.; and Kim, S. 2022. 3D GAN Inversion with Pose Optimization. arXiv:2210.07301.
- Ko, J.; Cho, K.; Choi, D.; Ryoo, K.; and Kim, S. 2023. 3d gan inversion with pose optimization. In *WACV*, 2967–2976.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. arXiv:2301.12597.
- Loshchilov, I.; and Hutter, F. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. arXiv:1608.03983.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2022. GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models. arXiv:2112.10741.
- Patashnik, O.; Wu, Z.; Shechtman, E.; Cohen-Or, D.; and Lischinski, D. 2021. Styleclip: Text-driven manipulation of stylegan imagery. In *ICCV*, 2085–2094.
- Poole, B.; Jain, A.; Barron, J. T.; and Mildenhall, B. 2022. DreamFusion: Text-to-3D using 2D Diffusion. arXiv:2209.14988.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763. PMLR.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 21(1): 5485–5551.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. arXiv:2204.06125.
- Roich, D.; Mokady, R.; Bermano, A. H.; and Cohen-Or, D. 2022. Pivotal tuning for latent-based editing of real images. *TOG*, 42(1): 1–13.

- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*, 10684–10695.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *NIPS*, 35: 36479–36494.
- Sanghi, A.; Chu, H.; Lambourne, J. G.; Wang, Y.; Cheng, C.-Y.; Fumero, M.; and Malekshan, K. R. 2022. Clip-forge: Towards zero-shot text-to-shape generation. In *CVPR*, 18603–18613.
- Shen, Y.; Gu, J.; Tang, X.; and Zhou, B. 2020. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 9243–9252.
- Song, J.; Meng, C.; and Ermon, S. 2022. Denoising Diffusion Implicit Models. arXiv:2010.02502.
- Tov, O.; Alaluf, Y.; Nitzan, Y.; Patashnik, O.; and Cohen-Or, D. 2021. Designing an encoder for stylegan image manipulation. *TOG*, 40(4): 1–14.
- Wang, K.; Wu, Q.; Song, L.; Yang, Z.; Wu, W.; Qian, C.; He, R.; Qiao, Y.; and Loy, C. C. 2020. Mead: A large-scale audio-visual dataset for emotional talking-face generation. In *ECCV*, 700–717. Springer.
- Wang, T.; Zhang, Y.; Fan, Y.; Wang, J.; and Chen, Q. 2022a. High-Fidelity GAN Inversion for Image Attribute Editing. In *CVPR*.
- Wang, T.; Zhang, Y.; Fan, Y.; Wang, J.; and Chen, Q. 2022b. High-fidelity gan inversion for image attribute editing. In *CVPR*, 11379–11388.
- Wei, T.; Chen, D.; Zhou, W.; Liao, J.; Tan, Z.; Yuan, L.; Zhang, W.; and Yu, N. 2022. Hairclip: Design your hair by text and reference image. In *CVPR*, 18072–18081.
- Wu, M.; Zhu, H.; Huang, L.; Zhuang, Y.; Lu, Y.; and Cao, X. 2023. High-Fidelity 3D Face Generation from Natural Language Descriptions. arXiv:2305.03302.
- Wu, Z.; Lischinski, D.; and Shechtman, E. 2021. Stylespace analysis: Disentangled controls for stylegan image generation. In *CVPR*, 12863–12872.
- Xia, W.; Yang, Y.; Xue, J.-H.; and Wu, B. 2021. Tedigan: Text-guided diverse face image generation and manipulation. In *CVPR*, 2256–2265.
- Xie, J.; Ouyang, H.; Piao, J.; Lei, C.; and Chen, Q. 2022. High-fidelity 3D GAN Inversion by Pseudo-multi-view Optimization. arXiv:2211.15662.
- Xu, T.; Zhang, P.; Huang, Q.; Zhang, H.; Gan, Z.; Huang, X.; and He, X. 2018. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *CVPR*, 1316–1324.
- Xu, Y.; Yang, Z.; and Yang, Y. 2023. SEEAAvatar: Photorealistic Text-to-3D Avatar Generation with Constrained Geometry and Appearance. arXiv:2312.08889.
- Yang, Y.; Zhuang, Y.; and Pan, Y. 2021. Multiple knowledge representation for big data artificial intelligence: framework, applications, and case studies. *Frontiers of Information Technology & Electronic Engineering*, 22(12): 1551–1558.
- Yang, Z.; Dong, J.; Liu, P.; Yang, Y.; and Yan, S. 2019. Very long natural scenery image prediction by outpainting. In *CVPR*, 10561–10570.
- Yin, F.; Zhang, Y.; Wang, X.; Wang, T.; Li, X.; Gong, Y.; Fan, Y.; Cun, X.; Shan, Y.; Oztireli, C.; and Yang, Y. 2023. 3D GAN Inversion with Facial Symmetry Prior. arXiv:2211.16927.
- Zhang, L.; Qiu, Q.; Lin, H.; Zhang, Q.; Shi, C.; Yang, W.; Shi, Y.; Yang, S.; Xu, L.; and Yu, J. 2023. DreamFace: Progressive Generation of Animatable 3D Faces under Text Guidance. arXiv:2304.03117.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 586–595.
- Zhou, D.; Yang, Z.; and Yang, Y. 2023. Pyramid Diffusion Models For Low-light Image Enhancement. In *IJCAI*.
- Zhou, T.; Wang, W.; Liang, Z.; and Shen, J. 2021. Face forensics in the wild. In *CVPR*.
- Zhu, J.; Shen, Y.; Zhao, D.; and Zhou, B. 2020. In-domain gan inversion for real image editing. In *ECCV*, 592–608. Springer.
- Zhu, J.-Y.; Krähenbühl, P.; Shechtman, E.; and Efros, A. A. 2016. Generative visual manipulation on the natural image manifold. In *ECCV*, 597–613. Springer.
- Zhu, Y.; Liu, H.; Song, Y.; Yuan, Z.; Han, X.; Yuan, C.; Chen, Q.; and Wang, J. 2022. One model to edit them all: Free-form text-driven image manipulation with semantic modulations. *NIPS*, 35: 25146–25159.