

Make Encoder Great Again in 3D GAN Inversion through Geometry and Occlusion-Aware Encoding

Ziyang Yuan^{1*} Yiming Zhu^{1,2*} Yu Li^{2†} Hongyu Liu³ Chun Yuan^{1†}

¹Tsinghua Shenzhen International Graduate School ²International Digital Economy Academy (IDEA)

³Hong Kong University of Science and Technology

<https://eg3d-goae.github.io>

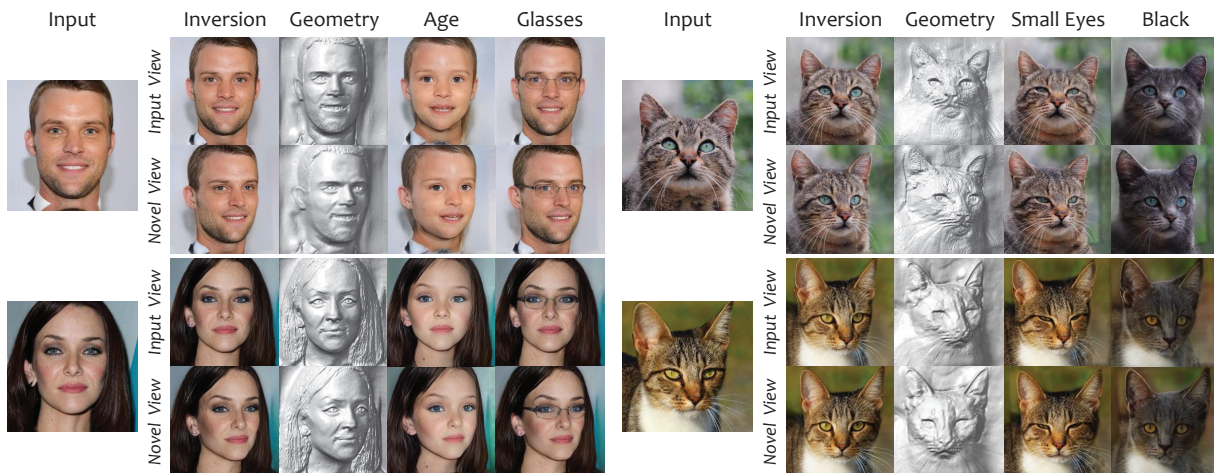


Figure 1: **The inversion and editing results of our method on human portraits and cat images.** The novel view synthesized images demonstrate that our method can achieve high-quality inversion and editing given a single image input.

Abstract

3D GAN inversion aims to achieve high reconstruction fidelity and reasonable 3D geometry simultaneously from a single image input. However, existing 3D GAN inversion methods rely on time-consuming optimization for each individual case. In this work, we introduce a novel encoder-based inversion framework based on EG3D, one of the most widely-used 3D GAN models. We leverage the inherent properties of EG3D’s latent space to design a discriminator and a background depth regularization. This enables us to train a geometry-aware encoder capable of converting the input image into corresponding latent code. Additionally, we explore the feature space of EG3D and develop an adaptive refinement stage that improves the representation ability of features in EG3D to enhance the recovery of fine-grained textural details. Finally, we propose an occlusion-aware fusion operation to prevent distortion in unobserved regions. Our method achieves impressive results comparable to optimization-based methods while operating up to 500 times faster. Our framework is well-suited for applications such as semantic editing.

* Equal contributions; † Corresponding author.

1. Introduction

Recent advances in 3D-aware generative adversarial networks (GAN) (e.g. [5, 6, 14, 32]) have integrated neural radiance fields (NeRF) [27] into style-based generation, resulting in remarkable success in generating highly realistic 3D portraits. Similar to 2D GAN inversion techniques, 3D GAN inversion also projects images into the latent space of pre-trained 3D GAN models. However, in addition to pursuing faithful image reconstruction and preserving fined details, 3D GAN inversion also aims for accurate 3D geometry, which is critical for generating realistic and consistent multi-view images. This ability to ensure the correctness of 3D geometry is a crucial advantage of 3D GAN inversion over its 2D counterpart.

Some optimization-based methods, such as [21, 47, 49], first invert the image to a latent code to obtain a coarse 3D shape. They then synthesize pseudo-multi-view images with this shape using a warping process. These pseudo images are used as supervision to optimize the parameters of the 3D GAN for more precise 3D geometry and better detail preservation. While these iterative optimization-based pipelines can prevent shape collapse and produce accurate

results, the optimization process needs to be performed for each individual case, which is time-consuming and impractical for many applications.

Several encoder-based methods [22, 23] leverage a pre-trained 3D GAN to generate proxy data by randomly sampling latent codes and camera poses. They then employ these paired proxy data to train an encoder in a self-supervised fashion and further refine features in 3D GAN. These encoder-based techniques offer higher efficiency but lower high-fidelity reconstruction ability. Moreover, due to a domain gap between the paired proxy data and real images, the encoder may invert the in-the-wild images to unsuitable latent codes. To address this issue, we study the properties of the latent space in 3D GAN and develop a geometry-aware encoder that predicts the latent code incorporating these properties. Furthermore, since the latent code can not preserve details of input images well, we explore the characteristics of feature space in 3D GAN to help our method improve the performance in reconstructing the original appearance. The latent and feature spaces are widely discussed in 2D GAN inversion [11, 24, 42].

This paper presents a novel encoder-based method for inverting 3D images generated by EG3D [5], a well-known and high-quality 3D GAN. Our approach is simple yet effective and is based on the discovery of a specific latent space within EG3D, which we call the **canonical latent space** \mathcal{W}_c . We find that latent codes belonging to this space can produce human faces with superior 3D shape and texture quality following two principles. Firstly, the input pose before the mapping network of EG3D must be a static canonical pose throughout the generation process. Secondly, the depth of the background should fall within a fixed range. Based on this, we propose using an encoder to invert the image into this canonical latent space using a proposed latent code discriminator and depth regularization. Although our encoder produces a suitable latent code that preserves the reasonable 3D shape of the input image, but the code is low dimension, making it challenging to capture texture details. To address this, similar to the 2D GAN inversion [3, 45], we refine the selected features in EG3D based on the residual between the input image and reconstruction image with predicted latent code. Meanwhile, we also find there is a **canonical feature space** \mathcal{F}_c in EG3D, so we design an alignment module to project the refined feature to the canonical feature space. Moreover, we introduce an occlusion-aware fusion operation to avoid texture and shape distortions in reconstructing invisible regions.

We conduct comprehensive experiments to evaluate the efficiency and effectiveness of our method, and the results demonstrate that our method achieves high-quality inversion comparable to optimization-based methods and superior to other encoder-based methods. We also demonstrate our approach is effective in 3D-aware editing applications.

See Fig. 1 for some examples of inversion and editing of human portraits and cat faces obtained by our method. Our contribution can be summarised as following:

- We conduct an exploration on the latent space and feature space of EG3D and discover the presence of a canonical attribute.
- Based on our analysis of the characteristics of the canonical latent space, we propose a geometry-aware encoder by utilizing a canonical latent discriminator and depth regularization.
- We introduce an alignment module based on the canonical feature space and an occlusion-aware method to refine the selected features of the generator and supplement high-quality details.
- The proposed method exhibits competitive performance with existing methods in human portraits, both qualitatively and quantitatively. Additionally, our method can also generalize to the cat faces and is effective in 3D editing tasks.

2. Related Work

2.1. Generative 3D-Aware Image Synthesis.

Generative Adversarial Network [13] made significant advancements in generating both 2D photorealistic images [19, 20] and 3D-Aware consistent images. In 3D-aware image generation, some works build 3D GAN based on explicit 3D representations, such as point-cloud [41], voxel [28, 29, 43], and mesh [12, 15, 17]. However, these models often fail to generate high-resolution 3D scenes with high-quality details.

Since NeRF [27] achieved outstanding multi-view image synthesis, a sequence of works has begun to use implicit 3D representation and volume rendering for 3D generation [5, 6, 10, 14, 30, 32]. Pi-GAN [6] is the first attempt at using GAN to generate NeRF for 3D-aware image synthesis. After that, StyleNeRF [14] proposes a view-consistent upsampling module to generate high-resolution images. StyleSDF [32] combines an SDF-based volume renderer and a 2D CNN network to obtain geometry information. GRAM [10] introduces 2D manifolds to regulate point sampling strategy in volume rendering. EG3D [5] leverages StyleGAN2 [20]’s structure and designs a computationally efficient 3D representation called tri-plane. It can synthesize high-resolution view-consistent images and produces high-quality 3D geometry. We explore EG3D’s latent space and build our inversion encoder based on EG3D.

2.2. GAN Inversion

GAN Inversion is an inverse process of the generation that projects the given image to corresponding latent code for faithful reconstruction from the source image. Thus,




Figure 2: **Recap of EG3D**. z, p denotes a Gaussian sample and a pose condition. The relation of w, F , and tri-plane is illustrated.

the image editing process can be conducted by semantic editing on the latent code [33, 37, 51]. In 2D GAN, inversion methods are mainly divided into three categories. Optimization-based methods [1, 8, 46] optimize the generator’s latent code or parameters both, achieve high reconstruction quality, but are time-consuming. Encoder-based methods [2, 35, 42] train an encoder to project input images to the latent codes. These methods can get faster inference than optimization-based methods but unsatisfactory inversion fidelity. Hybrid methods [4, 36, 51] use an encoder to get a latent as initialization, then optimize the latent code and generator parameters, which is still time-consuming.

Unlike 2D methods, 3D GAN Inversion should also consider multi-view-consistence. Directly transferring 2D methods to the 3D inversion is prone to suffering geometry collapse and artifacts. To solve it, Pose Opt. [21] jointly optimizes camera pose, latent codes, and generator parameters. HFGI3D [47] is an optimization-based method that extracts mesh to create pseudo-multi-view supervision. Meanwhile, encoder methods are also gradually developing. IDE-3D [40] uses semantic segmentation to re-train EG3D and learns an encoder for inversion to edit images. NeRF-3DE [23] uses identity contrastive learning to train a based encoder and then trains a refining encoder for better details. E3DGE [22] trains a base encoder to learn coarse shape and texture and supplements local details with pixel-aligned features. However, existing encoder-based methods [22, 23] rely on synthetic paired data and may not generalize well to real-world scenarios. None of the state-of-the-art methods can handle time efficiency and 3D inversion quality simultaneously.

3. EG3D Exploration

In this section, we first briefly review the EG3D model. Then, we explore and analyze the properties of latent space and feature space in EG3D.

The pipeline of EG3D is shown in Fig 2. It first combines a Gaussian vector z and a pose condition p as input and builds a mapper to generate a latent code w , this process can be denoted as $w = mapping(z, p)$. Then this latent code w will be sent to each layer of a style-based hierarchical gen-

erator to obtain the feature maps F . These feature maps are enlarged layer by layer. The last layer’s feature map reshape to *tri-plane*, an efficient 3D representation consists of three axis-aligned feature planes (F_{xy}, F_{xz}, F_{yz}). With this efficient 3D representation, EG3D samples point features from the tri-plane to synthesize multi-view images or depths by a rendering decoder R_{dec} with volume rendering [27]. Similar to the 2D GAN, the latent code w belongs to a latent space \mathcal{W} , and all feature maps in the generator construct the \mathcal{F} space.

3.1. Canonical Latent Space \mathcal{W}_c

In order to investigate the properties of the latent space \mathcal{W} and its relationship with pose conditioning, we adopt a fixed z and sample various pose conditions to generate corresponding latent code w . Fig. 3 illustrates that when w is conditioned on an unsuitable pose p far away from the canonical pose, the generated 3D shapes fail to generalize to other views synthesis. Only when the pose condition is canonical can the synthesized images maintain shape and texture consistency across multiple views. Meanwhile, we find a clear depth gap between the face and background regions when generating good results with the canonical pose condition. Under canonical pose, the depth of the background is restricted to a fixed range, while the face and background regions are confused for the other pose conditions.

These observations lead us to identify a specific subspace of \mathcal{W} intuitively, which we call **canonical latent space**, denoted as \mathcal{W}_c , that is responsible for ensuring reasonable 3D shape and texture across multiple views. There are two principles of this canonical latent space. Firstly, the input pose before the mapping network of EG3D must be a static canonical pose. Secondly, the depth of the background should fall within a fixed range to distinguish the face and background regions. Based on these findings, we propose that a high-quality 3D-GAN inversion method should first retrieve a w latent code in \mathcal{W}_c to ensure the best possible image synthesis results.

3.2. Canonical Feature Space \mathcal{F}_c

To analyze the characteristics of \mathcal{F} space, we generate the latent codes with different pose conditions to synthesize multiple tri-planes. We visualize the tri-plane as illustrated in Fig 3. The tri-plane also represents a canonical view when the pose condition is canonical (*i.e.*, the columns $g \sim i$ in row L_2). It indicates that there is also a specific subspace of \mathcal{F} , we call it **canonical feature space**, denoted as \mathcal{F}_c , and the canonical latent space \mathcal{W}_c produces the canonical feature space \mathcal{F}_c with the generator layers. The \mathcal{F}_c can supply sufficient facial information for rendering multi-view images.

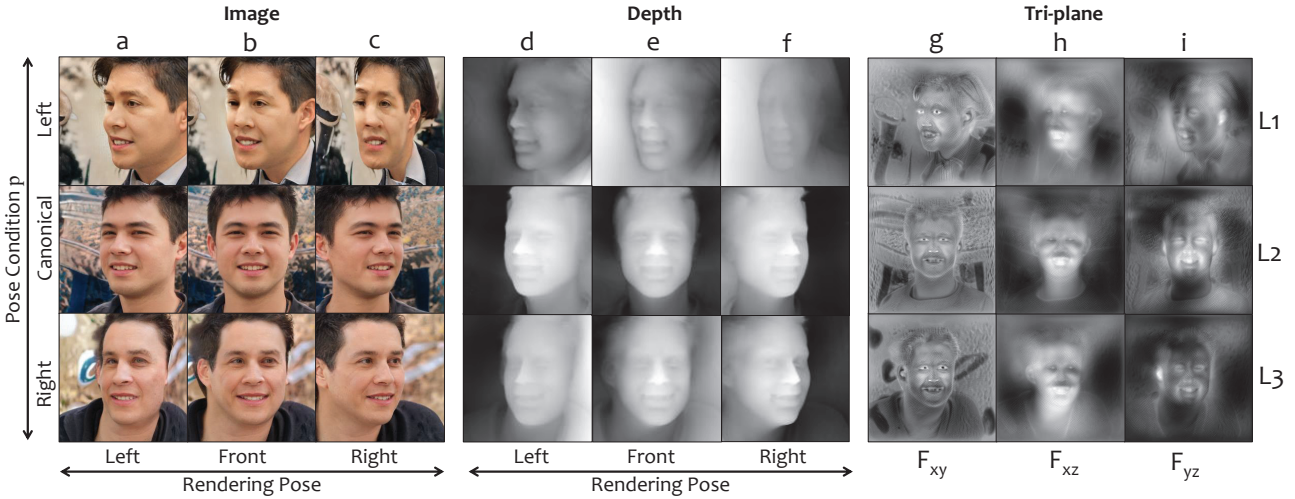


Figure 3: **Illustration of canonical Pose condition.** We synthesize multi-view images for three pose conditions p , including left, right, and canonical. As shown in $a \sim c$, the canonical pose condition in row L_2 has no distortion face shape. Moreover, depth maps in row L_2 demonstrate that images on the canonical pose condition have significant foreground and background differences. In contrast, flattened shape depth maps in rows L_1 and L_3 have foreground and background confusion. Columns $g \sim i$ reveal the tri-plane is also related to the pose condition.

4. Method

The pipeline of our method is shown in Fig. 4. We first project the input image I to the latent code $w^+ \subset \mathbb{R}^{14 \times 512}$ with our encoder, where $w^+ = (w_0, w_1, \dots, w_{13})$ is an extension of latent code w and can preserve the details better as mentioned in 2D GAN inversion [35, 42]. Then we calculate the residual between I and the generated image I_{w^+} based on w^+ to refine the features of the generator. In the following Sec. 4.1 and Sec. 4.2, we will describe how we ensure the latent code w^+ and refined feature belong to the canonical space, respectively. Furthermore, we introduce an occlusion-aware mix tri-plane in Sec. 4.3 to avoid the distortion of invisible regions reconstruction. In Sec. 4.4, we exhibit how to perform 3D semantic editing with our inversion method.

4.1. Geometry-aware Encoder

We design a geometry-aware encoder to obtain the latent code w^+ given the input image I and corresponding rendering camera pose c . We generate the inverted latent code w^+ according to the backbone’s pyramid feature progressively. It is worth noting that we choose Swin-transformer [25] as the backbone and further add attention modules at different scale feature layers for different level latent codes (more details of encoder architecture are in supplementary materials). After projecting the input image to the latent code w^+ , we can synthesize the $tri-plane_{w^+}$ through EG3D generator G and use rendering decoder R_{dec} to get the reconstruction image I_{w^+} under the rendering camera pose c .

$$\begin{aligned} w^+ &= E(I), \\ I_{w^+} &= R_{dec}(G(w^+), c) \end{aligned} \quad (1)$$

To make the inverted w^+ not deviate far from the canonical space \mathcal{W}_c discussed in Sec. 3.1, we use a canonical latent discriminator [31] to discriminate between $w_c \in \mathcal{W}_c$ and w^+ . The canonical latent discriminator, denoted as $\mathcal{D}_{\mathcal{W}_c}$, is trained in an adversarial scheme with the non-saturating GAN loss [13] and R_1 regularization [26], the loss of adversarial learning process can be written as below:

$$\begin{aligned} \mathcal{L}_{adv}^{\mathcal{D}_{\mathcal{W}_c}} &= -\mathbb{E}[\log \mathcal{D}_{\mathcal{W}_c}(w_c)] - \mathbb{E}[\log(1 - \mathcal{D}_{\mathcal{W}_c}(w_i))] \\ &\quad + \frac{\gamma}{2} \mathbb{E}[\|\nabla \mathcal{D}_{\mathcal{W}_c}(w_c)\|_2], \end{aligned} \quad (2)$$

$$\mathcal{L}_{adv}^E = -\mathbb{E}[\log \mathcal{D}_{\mathcal{W}_c}(w_i)], \quad (3)$$

where γ is the weight of R_1 regularization, the $\mathcal{L}_{adv}^{\mathcal{D}_{\mathcal{W}_c}}$ and \mathcal{L}_{adv}^E are the loss for the latent discriminator $\mathcal{D}_{\mathcal{W}_c}$ and encoder respectively. $\mathcal{D}_{\mathcal{W}_c}$ operates on w_c and each entry w_i of w^+ .

As discussed in Sec. 3.1, the background depth of an integrated face should be limited in a range. To estimate the background depth, we first random sample 0.1 million front images with the canonical pose condition, then we obtain the depth maps and background regions of these images. We use the parsing method [52] to get a binary mask M , where $M = 1$ indicates the background region and 0 otherwise. Finally, we calculate the average depth D_{avg} of background regions of these samples, and we set the D_{avg} as a regularization to constrain the background depth during inversion, the background regularization can be written as below:

$$\mathcal{L}_{BG} = \|D \odot M - D_{avg} \odot M\|_2. \quad (4)$$

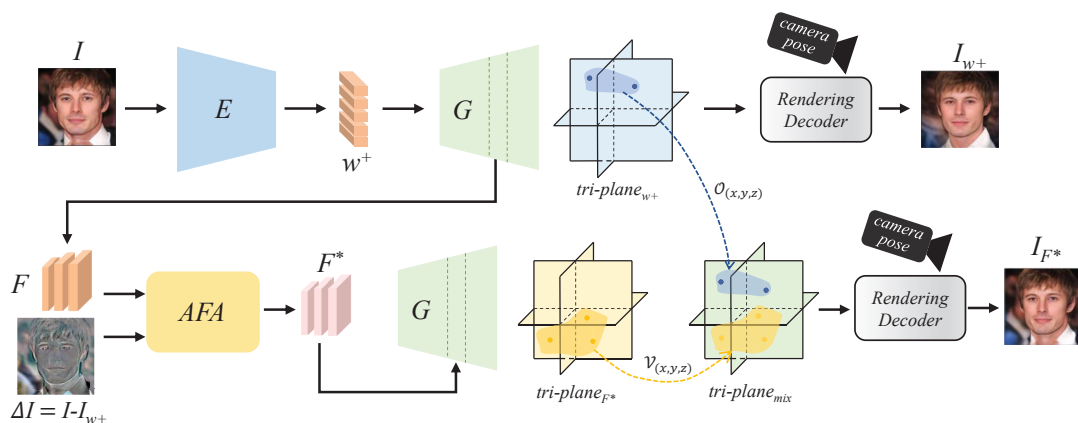


Figure 4: **Overview of our method.** Our framework could be divided into two parts. (1) \mathcal{W} space inversion. We design an encoder E to invert input image I into w^+ latent codes. The w^+ latent codes are fed into a pre-trained EG3D generator G to get $tri-plane_{w^+}$ and rendered into reconstruction image I_{w^+} . (2) Complement the \mathcal{F} space. We calculate the image residual ΔI between the input image and its reconstruction and propose AFA module to refine the F latent maps. The modified latent maps F^* are transformed into $tri-plane_{mix}$ by occlusion-aware mix and rendered into the fine detailed inversion image I_{F^*} .

Here, D and M are the depth map and background mask inverted image under the front-view pose.

In addition, we use reconstruction loss function \mathcal{L}_{rec} follow pSp [35] in the training stage. The \mathcal{L}_{rec} is composed by pixel-wise \mathcal{L}_2 loss, perceptual loss \mathcal{L}_{LPIPS} [50], and facial identity similarity loss \mathcal{L}_{ID} with pre-trained ArcFace network [9]:

$$\mathcal{L}_{rec} = \|I_{w^+} - I\|_2 + \mathcal{L}_{LPIPS}(I_{w^+}, I) + \mathcal{L}_{ID}(I_{w^+}, I). \quad (5)$$

4.2. Adaptive Feature Alignment

It is non-trivial to restore fine-detailed texture with only a highly compressed w latent code. Thus, adjusting high-dimension feature maps in the \mathcal{F} space is natural to achieve high-fidelity inversion as mentioned in related 2D GAN inversion methods [3, 11, 45]. Moreover, it is worth noting that the feature maps F generated with our inverted w^+ are in the canonical feature space \mathcal{F}_c .

We calculate the residual ΔI between I and I_{w^+} and hope to refine the feature maps to eliminate this residual. However, directly setting the residual as input and predicting the refined feature maps with a CNN to modify F is invalid. This is because the residual is highly entangled with the view of the input image, while F is in \mathcal{F}_c . When the view is not canonical, the refined feature maps with this residual can not project to the \mathcal{F}_c space simply. Thus, aligning the residual feature maps to the \mathcal{F}_c space is necessary to synthesize high-quality images, as mentioned in Sec. 3.2.

According to the above analysis, we propose an adaptive feature alignment (AFA) module that enables automatic canonical space alignment by cross-attention mechanism, as shown in Fig 5. First, we use a CNN to capture the feature $F_{\Delta I}$ of ΔI , and we take $F_{\Delta I}$ as the key K , value V , and the feature map F as a query Q . Then, we utilize cross-attention modules to align the $F_{\Delta I}$ with F and output a feature F_{align} . Since the F is already in the canonical feature




Figure 5: **The Architecture of AFA Module.** We use a CNN to capture the feature $F_{\Delta I}$ of residual image ΔI , and align the $F_{\Delta I}$ to F with cross attention module. Finally, we refine the F with the aligned feature F_{align} .

space \mathcal{F}_c , and the feature F_{align} can also have the properties of \mathcal{F}_c space. Finally, we use the F_{align} to predict the scale map γ and translation map β to refine the F with Feature-wise Linear Modulation (FiLM) [34]. The whole process of our AFA module can be formulated as follows:

$$\begin{aligned} F_{\Delta I} &= CNN(\Delta I), \\ Q, K, V &= W_Q F, W_K F_{\Delta I}, W_V F_{\Delta I}, \\ F_{align} &= Attention(Q, K, V), \\ \gamma, \beta &= Conv_{\gamma}(F_{align}), Conv_{\beta}(F_{align}), \\ F^* &= \gamma \odot F + \beta, \end{aligned} \quad (6)$$

where W_Q, W_K, W_V are learnable projection heads and d is the dimension of feature, the F^* is the final output of our AFA module. With this adaptive alignment, we achieve a detailed complement of \mathcal{F}_c space. Note that we adopt \mathcal{L}_{total} and a \mathcal{L}_2 regularization loss on $\Delta F = F^* - F$ when training the AFA module.

4.3. Occlusion-aware Mix Tri-plane

The image residual ΔI only contains information in the visible area of the input view. Using single-view residuals to modify the entire F will lead to distortion at the occlusion region. To solve this problem, we propose the Occlusion-

aware Mix Tri-plane.

The tri-plane contains all points in the 3D space which are divided into visible points set and occlusion points set according to the input view. Here, we define the 3D points viewed in the input image in tri-plane as the visible points set $\mathcal{V}_{(x,y,z)}$, points in the occlusion regions as occlusion points set $\mathcal{O}_{(x,y,z)}$. We can get the depth map D under the input view, and the two points set can be defined by:

$$\begin{aligned}\mathcal{V}_{(x,y,z)} &= \{[x, y, z]^T : [x, y, z]^T = \pi D(u, v) \mathcal{K}^{-1}[u, v, 1]^T\}, \\ \mathcal{O}_{(x,y,z)} &= \mathbb{R}^3 \setminus \mathcal{V}_{(x,y,z)},\end{aligned}\quad (7)$$

where $[u, v, 1]$ is the pixel homogeneous coordinate of the input image, π is the extrinsic camera parameter matrix of the input view, $D(u, v)$ is the depth of pixel $[u, v]$, \mathcal{K} is the intrinsic camera parameter matrix.

After determining the visible area, we can supplement the refined details of the input view to the tri-plane. We define the occlusion-aware mix tri-plane as follows:

$$\begin{aligned}\text{tri-plane}_{w^+} &= G(w^+), \\ \text{tri-plane}_{F^*} &= G(F^*), \\ \text{tri-plane}_{mix} &= \text{tri-plane}_{F^*}|_{(x,y,z) \in \mathcal{V}} \\ &\quad + \text{tri-plane}_{w^+}|_{(x,y,z) \in \mathcal{O}},\end{aligned}\quad (8)$$

where w^+ is the latent code learned from the encoder, and F^* is the refined feature map. Finally, we use tri-plane_{mix} with rendering decoder to synthesize images under different views when inference.

4.4. Editing

To achieve 3D-consistent editing, we need to modify both w^+ and F^* . We first get the edited latent code \hat{w}^+ through classic latent code editing methods [33, 37]. Then for F^* , we follow FS [48] to extract the features $F^{\hat{w}^+}$ and F^w from the generator convolution layer with input \hat{w}^+ and w , respectively. Finally, we calculate the difference between these two features and add it to the F^* to get the edited features \hat{F}^* :

$$\hat{F}^* = F^* + F^{\hat{w}^+} - F^w. \quad (9)$$

With the modified \hat{w}^+ and \hat{F}^* , we use the method in section 4.3 to get edited mix tri-plane, and use rendering decoder to synthesize 3D-consistent edited images.

4.5. Losses.

In the first training stage, we use the loss function mentioned in Sec 4.1. For the input view reconstruction, the loss \mathcal{L}_{rec} is computed by:

$$\mathcal{L}_{rec} = \lambda_1 \|I_{w^+} - I\|_2 + \lambda_2 \mathcal{L}_{LPIPS}(I_{w^+}, I) + \lambda_3 \mathcal{L}_{ID}(I_{w^+}, I), \quad (10)$$

where $\lambda_1 = 1.0, \lambda_2 = 0.8, \lambda_3 = 0.25$. The total loss for training our encoder and the canonical latent discriminator can be written as:

$$\mathcal{L}_{total}^E = \mathcal{L}_{adv}^{D_{w^c}} + \lambda_4 \mathcal{L}_{adv}^E + \lambda_5 \mathcal{L}_{BG} + \mathcal{L}_{rec}, \quad (11)$$

where R_1 regularization weight $\gamma = 10, \lambda_4 = 0.05, \lambda_5 = 5.0$. Moreover, we apply regularization on $\Delta w^+ = (w_1 - w_0, w_2 - w_0, \dots, w_{13} - w_0)$ by \mathcal{L}_2 norm, i.e. $\mathcal{L}_{w^+} = \lambda_6 \|\Delta w^+\|_2$, where $\lambda_6 = 0.001$.

Once the loss in the first stage is converged, we freeze the encoder parameters and drop the canonical latent discriminator and losses $\mathcal{L}_{adv}^{D_{w^c}}, \mathcal{L}_{adv}^E, \mathcal{L}_{w^+}$, but add a \mathcal{L}_2 regularization on $\Delta F = F^* - F$ with $\lambda_7 = 0.0001$ for training the AFA module, i.e. $\mathcal{L}_{\Delta F} = \lambda_7 \|\Delta F\|_2$. We apply the loss settings to both human data and cat data.

For re-implemented pSp [35] and e4e [42], we train it with the same loss weight in our settings, and train them until the losses converge.

5. Experiments

Experiment setup. We evaluate our method on human faces and cat faces. For human faces, we use FFHQ [1] dataset for training, and CelebA-HQ [18] and MEAD [44] for evaluation. For the cat faces, we use AFHQ-CAT [7] for training and evaluation. In addition, we employ the pre-trained EG3D generator and obtain the corresponding camera parameters of each image following the official implementation in EG3D.

Baselines. We compare our method with encoder (IDE-3D [40]) and optimization-based (PTI [36] and Pose Opt. [21]) baselines in these two aspects. Moreover, we build two encoder baselines imitating the state-of-the-art StyleGAN inverters, pSp [35] and e4e [42], and present the model structure in the supplementary.

5.1. Qualitative analysis

Source View Inversion Results. Fig. 6 shows the inversion performance comparison between ours and other methods. Our method’s hair details and facial similarity are obviously competitive with optimization methods and surpass all the encoder-based methods.

Novel View Synthesis Results. The comparison results are shown in Fig. 6. Inspecting the eye gaze direction and person’s identity preservation when changing to novel views can reveal the differences. It can be seen that on some extreme side-face camera poses, some optimization methods produce artifacts and geometry distortions. Our method can keep synthesizing 3D-consistent images faithful to the person’s identity in the input.

5.2. Quantitative analysis

We calculate the mean square error (MSE), LPIPS [50], and FID [16] between the source image and its inversion.

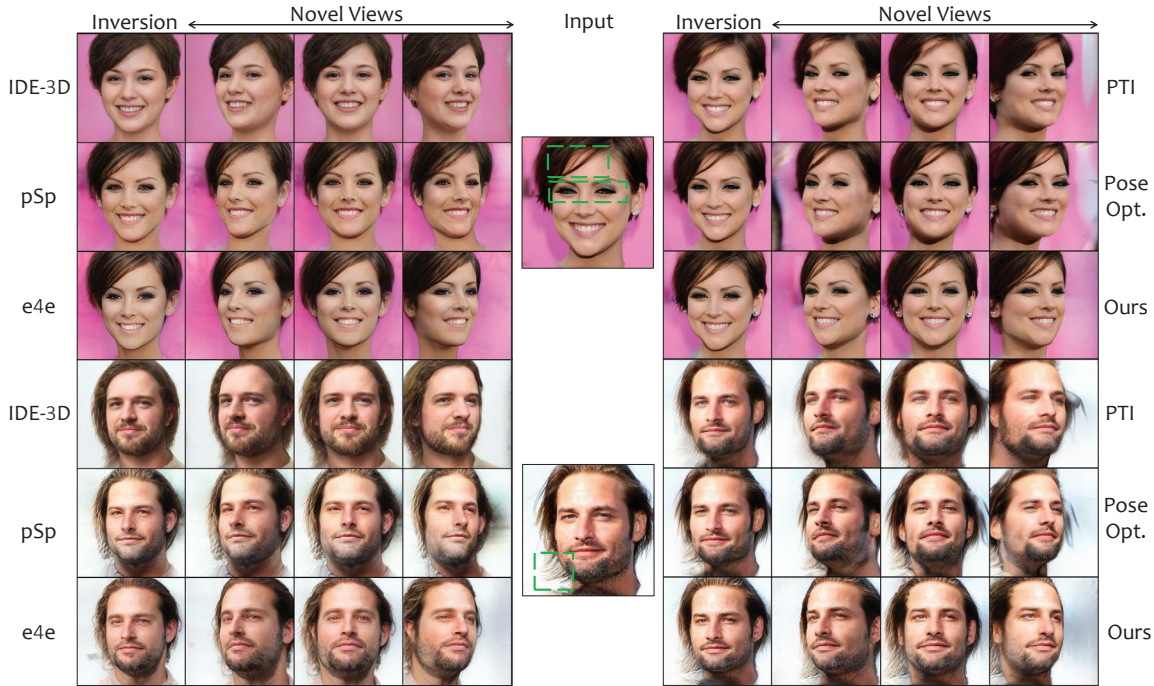


Figure 6: **Visual comparison of inversion and novel-view synthesized results between our method and others.** Our method surpasses all the encoder-based methods and is competitive with optimization methods in recovering texture details, *i.e.* our method prevents the eye orientation from changing with the camera view and artifacts in the novel view.

Category	Method	Reconstruction Quality						Novel View Quality							
		Input View						yaw=±60°				yaw=0°			
		MSE ↓	LPIPS ↓	FID ↓	ID ↑	Geo. Err. ↓	Time (s) ↓	LPIPS ↓	FID ↓	ID ↑	Geo. Err. ↓	LPIPS ↓	FID ↓	ID ↑	Geo. Err. ↓
Optimization	PTI [36]	0.0163	0.0918	18.0	0.8417	0.0850	115.2	0.2926	69.8	0.621	0.1481	0.1925	53.5	0.718	0.1006
	Pose Opt. [21]	0.0164	0.0879	17.5	0.8390	0.0759	219.7	0.4477	132	0.321	0.1763	0.2436	60.2	0.605	0.1504
Encoder-based	IDE-3D [40]	0.1056	0.2806	43.6	0.7194	0.1835	0.0659	0.3317	79.0	0.505	0.2914	0.3074	63.7	0.513	0.2528
	pSp-3D [35]	0.0726	0.2082	31.2	0.7817	0.1846	0.1284	0.3223	65.7	0.591	0.1287	0.3132	65.2	0.671	0.1789
	e4e-3D [42]	0.0808	0.2281	33.0	0.7359	0.1161	0.1303	0.2495	57.7	0.588	0.1269	0.2394	53.8	0.653	0.1305
	Ours- w^+	0.0697	0.1981	30.9	0.7873	0.1100	0.0963	0.2303	52.2	0.655	0.0943	0.2372	54.1	0.685	0.0955
	Ours	0.0301	0.1257	20.5	0.8752	0.0987	0.2262	0.2192	50.9	0.700	0.0950	0.2204	51.7	0.743	0.0944

Table 1: **Comparison with state-of-the-art methods on inversion performance.** The numbers in bold are the best performance of the optimization and encoder-based methods, respectively. Metrics of Reconstruction Quality are calculated under the input view. Metrics of Novel View Quality are the average of multi-view. The yaw angle is the source view angle.

We adopt a pre-trained face key point detector parsenet [52] to calculate the identity similarity. To measure the geometry properties, we followed [5, 38], utilizing a pre-trained 3D face reconstruction model to extract a “pseudo” ground truth depth map from the source image. Then, we obtain the depth of inversion and normalize all depth maps to zero means and unit variance to calculate geometry error (Geo. Err.) by using average \mathcal{L}_2 distance.

Source View Inversion Results. Table 1 presents the quantitative comparison of the input view image reconstruction on CelebA-HQ test dataset. Our method is two orders of magnitude faster than the optimized methods with similar

performance, which is notably better than other encoder-based methods.

Novel View Synthesis Results. In order to measure the multi-view consistency and the synthesized quality, we evaluate our method and baselines with a multi-view dataset, MEAD [44]. Each person in MEAD has five images with increasing yaw angles. We select one angle as the input view, invert the input image and synthesize the other four novel views. Then we use the same metrics to calculate the difference between the synthesized images and the ground truth. The quantitative comparison between the novel view synthesized results and the ground truth in dif-



Figure 7: **Ablation of Geometry-aware Encoder.** All images are synthesized only with w^+ . Without both modules, w^+ fails to generate integrated 3D geometry and texture. The L_{BG} helps distinguish foreground and background. With \mathcal{D}_{W_c} , we can get good inversion under the input view but fail under the novel view. With both modules, w^+ can obtain an integrated side face of the woman, which is occlusion under the input view.

ferent views is shown in Table 1. Our method shows similar performance with PTI when the input image is a front face (yaw = 0°). When inverting side faces, the difficulty of inversion increases drastically when the yaw angle increases. We calculated the metrics of the largest yaw angle in MEAD (yaw = ±60°), and our model has robust performance when the test yaw angle varies. However, the performance of optimization methods degrades remarkably at an extreme angle.

5.3. Ablation Analysis

We present ablation studies with qualitative analysis of our proposed method below, and the corresponding numerical ablation results are in our supplementary.

Geometry-aware Encoder. We conduct ablation experiments and make qualitative comparisons to verify that our proposed canonical latent discriminator (\mathcal{D}_{W_c}) and background regularization (\mathcal{L}_{BG}) contribute to restoring 3D geometry. Generally, recovering the geometry from a side face input is more challenging and we test such a case. As shown in Fig. 7, only if we adopt the canonical discriminator and the background regularization could we obtain a satisfactory foreground and background geometry.

Adaptive Feature Alignment. We proposed Adaptive Feature Alignment (AFA) to refine the F latent map and add texture details from the source view. Fig. 8 demonstrates that we can recover more details with AFA.

Occlusion-aware Mix Tri-plane. To verify the necessity of our proposed occlusion-aware mix tri-plane, we ablate this operation and visualize the novel view results in Fig. 9. The occlusion map shows the spatial position of surface points



Figure 8: **Ablation of Adaptive Feature Alignment.** We replace AFA with a CNN network and obtain the results. Without feature alignment, it is difficult to reconstruct the hair and facial details.




Figure 9: **Ablation of Occlusion-aware mix tri-plane.** We visualize the occlusion regions $\mathcal{O}_{(x,y,z)}$ of the input view in the last column (black points). The green arrow points out $\mathcal{O}_{(x,y,z)}$ in the novel view, and it is obvious that the occlusion-aware mix tri-plane can avoid artifacts.

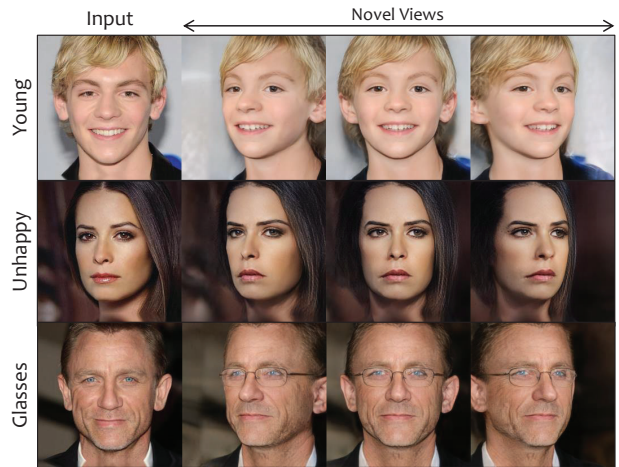


Figure 10: **Editing result under novel views.** We add three different attributes to three faces. Our method generates both realistic editing and 3D-consistent novel views.

in ray sampling, the visible region in the original view is white, and the occlusion region is black. Without the mix tri-plane, the occlusion area from the source view is easy to produce artifacts in novel views.

5.4. Image Editing Results

In GAN inversion-based image manipulation, the editability and the reconstruction trade-off widely exists in both 2D and 3D scenarios. We validate the editing method




Figure 11: Results of our encoder in other 3D GANs.

	MSE ↓	LPIPS ↓	FID ↓		MSE ↓	LPIPS ↓	FID ↓
w/o AFA	0.0586	0.2317	13.43	w/ AFA	0.0367	0.1745	10.82

Table 2: Ablation of AFA on ShapeNet Cars.

mentioned in Sec. 4.4 and display the multi-view edited results in Fig. 10. We choose three different edit directions, 'Young' to reduce person's age, 'Unhappy' to change emotion, and 'Glasses' to add accessories. It is noticeable that our method not only restores the geometric and texture information of the input image but also ensures the editing consistency of multi-view. The edited results demonstrate that our inversion method can be well applied to the facial image editing tasks.

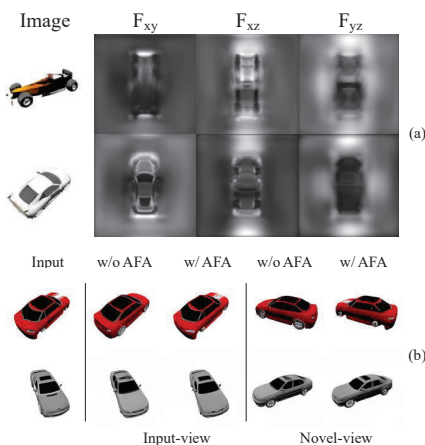


Figure 12: (a) Tri-plane canonical feature space. (b) Results of our method on ShapeNet Cars datasets. Best viewed zoomed in and in color.

6. Generalization

Our method can be adapted to other 3D GANs (e.g., EpiGRAF [39] and StyleSDF [32]) and it works also with categories where pose estimation and canonicalization are more difficult (e.g., Shapenet cars).

Extension to other 3D GANs. We choose EG3D as our base 3D GAN due to its superior performance, but our method can also be applied to other 3D GANs. The canon-

ical space, which plays a vital role in 3D GAN inversion, is commonly presented in the latent space (*i.e.*, w^+ space) of 3D GANs. By inverting images into the canonical space, we preserve geometry and achieve high-fidelity reconstruction, as demonstrated in our experiments on EpiGRAF [39] and StyleSDF [32] in Fig. 11. The results of w/o D_{w_c}, L_{BG} mean that we only use the base encoder without the canonical latent discriminator and the background depth regularization. In contrast, w/ D_{w_c}, L_{BG} is our full encoder model. The noticeable performance gap proves the powerful inversion ability of our method. Our method explores the canonical space's characteristics, ensuring distortion-free shape generation and exceptional performance in image 3D inversion.

Extension to ShapeNet Cars. Fig. 12 shows (a) the canonical space and (b) the inversion results on ShapeNet Cars. Even without identity loss, our method can achieve good performance (see the Col.3,5 of (b)). As the tri-plane is canonical, without the AFA module, the residual feature can not align with the canonical feature space (see the Col. 2,4 of (b) and Table 2).

7. Conclusion

In this work, we explore the inherent property of EG3D's latent space and find that the key to achieve high-quality shape and texture inversion is inverting input images to the canonical space. In light of this, we design an efficient geometry-aware encoder with an adaptive feature alignment module and demonstrate their effectiveness. Besides, we propose an occlusion-aware mix tri-plane to avoid invisible region distortion in novel views. Extensive experiments demonstrate that our method has competitive inversion performance with optimization-based methods, but operates two orders of magnitude faster. In addition, our method is demonstrated to be a powerful tool for realistic and 3D view-consistent editing of facial images.

8. Acknowledgement

This work was supported by the National Key R&D Program of China (2022YFB4701400/4701402), SZSTC Grant (JCYJ20190809172201639, WDZC20200820200655001), Shenzhen Key Laboratory (ZDSYS20210623092001004), and Beijing Key Lab of Networked Multimedia.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4432–4441, 2019. 3, 6
- [2] 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. 3
- [3] Yuval Alaluf, Omer Tov, Ron Mokady, Rinon Gal, and Amit Bermano. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18511–18521, 2022. 2, 5
- [4] Lucy Chai, Jun-Yan Zhu, Eli Shechtman, Phillip Isola, and Richard Zhang. Ensembling with deep generative views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14997–15007, 2021. 3
- [5] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16123–16133, 2022. 1, 2, 7
- [6] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5799–5809, 2021. 1, 2
- [7] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2020. 6
- [8] Edo Collins, Raja Bala, Bob Price, and Sabine Susstrunk. Editing in style: Uncovering the local semantics of gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5771–5780, 2020. 3
- [9] 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
- [10] Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. Gram: Generative radiance manifolds for 3d-aware image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10673–10683, 2022. 2
- [11] Tan M Dinh, Anh Tuan Tran, Rang Nguyen, and Binh-Son Hua. Hyperinverter: Improving stylegan inversion via hypernetwork. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11389–11398, 2022. 2, 5
- [12] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *European Conference on Computer Vision*, 2020. 2
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014. 2, 4
- [14] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985*, 2021. 1, 2
- [15] Paul Henderson, Vagia Tsiminaki, and Christoph H Lampert. Leveraging 2d data to learn textured 3d mesh generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems*, 30, 2017. 6
- [17] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *European Conference on Computer Vision*, 2018. 2
- [18] 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. 6
- [19] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [20] 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 Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [21] Jaehoon Ko, Kyusun Cho, Daewon Choi, Kwangrok Ryoo, and Seungryong Kim. 3d gan inversion with pose optimization. *arXiv preprint arXiv:2210.07301*, 2022. 1, 3, 6, 7
- [22] Yushi Lan, Xuyi Meng, Shuai Yang, Chen Change Loy, and Bo Dai. Self-supervised geometry-aware encoder for style-based 3d gan inversion. *arXiv preprint arXiv:2212.07409*, 2022. 2, 3
- [23] Yu-Jhe Li, Tao Xu, Bichen Wu, Ningyuan Zheng, Xiaoliang Dai, Albert Pumarola, Peizhao Zhang, Peter Vajda, and Kris Kitani. 3d-aware encoding for style-based neural radiance fields. *arXiv preprint arXiv:2211.06583*, 2022. 2, 3
- [24] Hongyu Liu, Yibing Song, and Qifeng Chen. Delving stylegan inversion for image editing: A foundation latent space viewpoint. *arXiv preprint arXiv:2211.11448*, 2022. 2
- [25] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 4
- [26] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International Conference on Machine Learning*, pages 3481–3490. PMLR, 2018. 4
- [27] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf:

- Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, 2020. 1, 2, 3
- [28] 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. 2
- [29] Thu H Nguyen-Phuoc, Christian Richardt, Long Mai, Yongliang Yang, and Niloy Mitra. Blockgan: Learning 3d object-aware scene representations from unlabelled images. *Advances in Neural Information Processing Systems*, 33:6767–6778, 2020. 2
- [30] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021. 2
- [31] Yotam Nitzan, Amit Bermano, Yangyan Li, and Daniel Cohen-Or. Face identity disentanglement via latent space mapping. *arXiv preprint arXiv:2005.07728*, 2020. 4
- [32] 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, 9
- [33] 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. 3, 6
- [34] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 5
- [35] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2287–2296, 2021. 3, 4, 5, 6, 7
- [36] 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. 3, 6, 7
- [37] Yujun Shen, Jinjin Gu, Xiaou 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. 3, 6
- [38] Yichun Shi, Divyansh Aggarwal, and Anil K Jain. Lifting 2d stylegan for 3d-aware face generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6258–6266, 2021. 7
- [39] Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. Epigraf: Rethinking training of 3d gans. *arXiv preprint arXiv:2206.10535*, 2022. 9
- [40] Jingxiang Sun, Xuan Wang, Yichun Shi, Lizhen Wang, Jue Wang, and Yebin Liu. Ide-3d: Interactive disentangled editing for high-resolution 3d-aware portrait synthesis. *arXiv preprint arXiv:2205.15517*, 2022. 3, 6, 7
- [41] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Multi-view 3D models from single images with a convolutional network. In *European Conference on Computer Vision*, 2016. 2
- [42] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 2, 3, 4, 6, 7
- [43] Shubham Tulsiani, Hao Su, Leonidas J. Guibas, Alexei A. Efros, and Jitendra Malik. Learning shape abstractions by assembling volumetric primitives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [44] Kaisiyuan Wang, Qianyi Wu, Linsen Song, Zhuoqian Yang, Wayne Wu, Chen Qian, Ran He, Yu Qiao, and Chen Change Loy. Mead: A large-scale audio-visual dataset for emotional talking-face generation. In *European Conference on Computer Vision*, pages 700–717. Springer, 2020. 6, 7
- [45] Tengfei Wang, Yong Zhang, Yanbo Fan, Jue Wang, and Qifeng Chen. High-fidelity gan inversion for image attribute editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11379–11388, 2022. 2, 5
- [46] Zongze Wu, Dani Lischinski, and Eli Shechtman. Stylespace analysis: Disentangled controls for stylegan image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12863–12872, 2021. 3
- [47] Jiaxin Xie, Hao Ouyang, Jingtian Piao, Chenyang Lei, and Qifeng Chen. High-fidelity 3d gan inversion by pseudo-multi-view optimization. *arXiv preprint arXiv:2211.15662*, 2022. 1, 3
- [48] Xu Yao, Alasdair Newson, Yann Gousseau, and Pierre Hellier. A style-based gan encoder for high fidelity reconstruction of images and videos. *European Conference on Computer Vision*, 2022. 6
- [49] Fei Yin, Yong Zhang, Xuan Wang, Tengfei Wang, Xiaoyu Li, Yuan Gong, Yanbo Fan, Xiaodong Cun, Ying Shan, Cengiz Oztireli, et al. 3d gan inversion with facial symmetry prior. *arXiv preprint arXiv:2211.16927*, 2022. 1
- [50] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 5, 6
- [51] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *European Conference on Computer Vision*, pages 592–608. Springer, 2020. 3
- [52] zllrunning. face-parsing.pytorch. <https://github.com/zllrunning/face-parsing.PyTorch/>, 2019. 4, 7