

Self-Supervised Geometry-Aware Encoder for Style-Based 3D GAN Inversion

Yushi Lan¹ Xuyi Meng¹ Shuai Yang¹ Chen Change Loy¹ Bo Dai²
¹S-Lab, Nanyang Technological University, Singapore ²Shanghai AI Laboratory

Abstract

StyleGAN has achieved great progress in 2D face reconstruction and semantic editing via image inversion and latent editing. While studies over extending 2D StyleGAN to 3D faces have emerged, a corresponding generic 3D GAN inversion framework is still missing, limiting the applications of 3D face reconstruction and semantic editing. In this paper, we study the challenging problem of 3D GAN inversion where a latent code is predicted given a single face image to faithfully recover its 3D shapes and detailed textures. The problem is ill-posed: innumerable compositions of shape and texture could be rendered to the current image. Furthermore, with the limited capacity of a global latent code, 2D inversion methods cannot preserve faithful shape and texture at the same time when applied to 3D models. To solve this problem, we devise an effective self-training scheme to constrain the learning of inversion. The learning is done efficiently without any real-world 2D-3D training pairs but proxy samples generated from a 3D GAN. In addition, apart from a global latent code that captures the coarse shape and texture information, we augment the generation network with a local branch, where pixel-aligned features are added to faithfully reconstruct face details. We further consider a new pipeline to perform 3D view-consistent editing. Extensive experiments show that our method outperforms state-of-the-art inversion methods in both shape and texture reconstruction quality. Code and data will be released.

synthesize 3D view-consistent editing results, *e.g.*, changing a neutral expression to smiling, by altering the estimated latent code.

GAN inversion [49] has been extensively studied for 2D images but remains underexplored in the 3D world. Inversion can be achieved via optimization [1, 2, 38], which typically provides a precise image-to-latent mapping but can be time-consuming, or encoder-based techniques [37, 44, 47], which explicitly learn an encoding network that maps an image into the latent space. Encoder-based techniques enjoy faster inversion, but the mapping is typically inferior to optimization. In this study, we extend the notion of encoder-based inversion from 2D images to 3D shapes.

Adding the additional dimension makes inversion more challenging beyond the goal of reconstructing an editable shape with detail preservation. In particular, **1)** Recovering 3D shapes from 2D images is an ill-posed problem, where innumerable compositions of shape and texture could generate identical rendering results. 3D supervisions are crucial to alleviate the ambiguity of shape inversion from images. Though high-quality 2D datasets are easily accessible, owing to the expensive cost of scans there is currently a lack of large-scale labeled 3D datasets. **2)** The global latent code, due to its compact and low-dimensional nature, only captures the coarse shape and texture information. Without high-frequency spatial details, we cannot generate high-fidelity outputs. **3)** Compared with 2D inversion methods where the editing view mostly aligns with the *input view*, in 3D editing we expect the editing results to perform well over the *novel views* with large pose variations. Therefore, 3D GAN inversion is non-trivial task and could not be achieved by directly applying existing approaches.

To this end, we propose a novel Encoder-based **3D GAN invErsion** framework, E3DGE, which addresses the aforementioned three challenges. Our framework has three novel components with a delicate model design. Specifically:

Learning Inversion with Self-supervised Learning - The first component focuses on the training of the inversion encoder. To address the shape collapse of single-view 3D reconstruction without external 3D datasets, we retrofit the generator of a 3D GAN model to provide us with diverse pseudo training samples, which can then be used to train

1. Introduction

The main goal of this work is to devise an effective approach for encoder-based 3D Generative Adversarial Network (GAN) inversion. In particular, we focus on the reconstruction of 3D face, requiring just a single 2D face image as the input. In the inversion process, we wish to map a given image to the latent space and obtain an editable latent code with an encoder. The latent code will be further fed to a generator to reconstruct the corresponding 3D shape with high-quality shape and texture. Further to the learning of an inversion encoder, we also wish to develop an approach to

our inversion encoder in a self-supervised manner. Specifically, we generate 3D shapes from the latent space \mathcal{W} of a 3D GAN, and then render diverse 2D views from each 3D shape given different camera poses. In this way, we can generate many pseudo 2D-3D pairs together with the corresponding latent codes. Since the pseudo pairs are generated from a smooth latent space that learns to approximate a natural shape manifold, they serve as effective surrogate data to train the encoder, avoiding potential shape collapse.

Local Features for High-Fidelity Inversion - The second component learns to reconstruct accurate texture details. Our novelty here is to leverage local features to enhance the representation capacity, beyond just the global latent code generated by the inversion encoder. Specifically, in addition to inferring an editable global latent code to represent the overall shape of the face, we further devise an hour-glass model to extract local features over the residuals details that the global latent code fails to capture. The local features, with proper projection to the 3D space, serve as conditions to modulate the 2D image rendering. Through this effective learning scheme, we marry the benefits of both global and local priors and achieve high-fidelity reconstruction.

Synthesizing View-consistent Edited Output - The third component addresses the problem of novel view synthesis, a problem unique to 3D shape editing. Specifically, though we achieve high-fidelity reconstruction through aforementioned designs, the local residual features may not fully align with the scene when being semantically edited. Moreover, the occlusion issue further degrades the fusion performance when rendering from novel views with large pose variations. To this end, we propose a 2D-3D hybrid alignment module for high-quality editing. Specifically, a 2D alignment module and a 3D projection scheme are introduced to jointly align the local features with edited images and inpaint occluded local features in novel view synthesis.

Extensive experiments show that our method achieves 3D GAN inversion with plausible shapes and high-fidelity image reconstruction without affecting editability. Owing to the self-supervised training strategy with delicate global-local design, our approach performs well on real-world 2D and 3D benchmarks without resorting to any real-world 3D dataset for training. To summarize, our main contributions are as follows:

- We propose an early attempt at learning an encoder-based 3D GAN inversion framework for high-quality shape and texture inversion. We show that, with careful design, samples synthesized by a GAN could serve as proxy data for self-supervised training in inversion.
- We present an effective framework that uses local features to complement the global latent code for high-fidelity inversion.
- We propose an effective approach to synthesize view-

consistent output through a 2D-3D hybrid alignment module.

2. Related Work

3D-aware Image Synthesis. Generative Adversarial Network [13] has shown promising results in generating photorealistic images [6, 21, 22] and inspired researchers to put efforts on 3D aware generation [15, 29, 32]. However, these methods use explicit shape representations, *i.e.*, voxels [15, 29] and meshes [32] as the intermediate shape models, which lacks photorealism and is memory-inefficient. Motivated by the recent success of neural rendering [26, 27, 35], researchers shift to implicit function along with the volume rendering process as the incorporated 3D inductive bias. Among them, NeRF [27] proposed an implicit 3D representation for novel view synthesis which defines a scene as $\{c, \sigma\} = F_{\Phi}(x, v)$, where x is the query point, v is the viewing direction from camera origin to x , c is the emitted radiance (RGB value), σ is the volume density. Researchers further extend NeRF to generation task [7, 42] and show impressive view-consistency on the synthesized results. To increase the generation resolution, recent works [8, 16, 51] resort to voxel-based representations or adopting a hybrid design [8, 14, 30, 31]. By lifting the intermediate low-resolution 2D features to high resolution with a 2D super-resolution decoder, the hybrid design achieves view-consistent synthesis at high resolution, *e.g.*, 1024^2 . Beyond synthesizing realistic and diverse images, previous works [4, 17, 18, 34, 52, 54] have shown that pre-trained generators of GAN can be viewed as a compressed and organized training dataset. Through careful design in the sampling strategy [18], loss functions [34] and generation process [54], off-the-shelf image generators could facilitate a series of downstream visual applications.

2D GAN Inversion. To leverage the strong priors encoded in GANs, GAN inversion techniques on 2D GANs are well developed. Optimization-based methods [1, 2] could achieve photorealistic reconstruction at the cost of slow inference and lack of editability. Encoder-based methods [9, 37, 44, 47, 55] have been developed to speed up the inversion and show better properties in editing through specific model design [37, 47] and training strategies [44]. pSp [37] proposed an encoder architecture designed for human faces, serving as the backbone for many approaches. e4e [44] analyzed the trade-offs between editability and fidelity. However, they [1, 2, 37, 44, 55] all adopt global latent code alone for GAN inversion task, thus failing to recover high-fidelity details. Recently, HFGI [47] introduce an extra spatial consultation map to mitigate this issue, though still designed to restore 2D textures without considering 3D shape modeling. In this work, we propose a delicate design that exploits local features to recover texture details and achieves view-consistent synthesis.

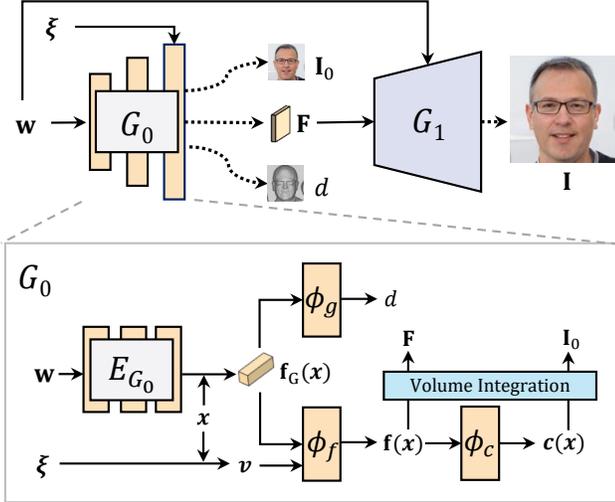


Figure 1. **StyleSDF**. Given a sampled latent code w and a camera pose ξ , StyleSDF generates object SDF d to depict the shape and the corresponding face image I .

3. Preliminaries

Hybrid 3D-aware Generation. To achieve high-resolution novel view synthesis, hybrid 3D-aware generator [8, 14, 30, 31] is proposed. It is a cascade model $G = G_0 \circ G_1$ composed of a NeRF-based renderer G_0 [7] and a 2D super-resolution network G_1 , as shown in Fig. 1. Both G_0 and G_1 follow the style-based architecture [21, 23] to accept a latent code w to control the style of the generated object. During generation, G_0 captures the underlying geometry with the full control of w and camera pose ξ , and renders a low resolution image I_0 and an intermediate feature map F . Then, G_1 further upsamples F to obtain a high-resolution image I with added high-frequency details.

Among them, StyleSDF [31] introduces signed distance function (SDF) to serve as a proxy for the density function $\sigma(x)$ used for the volume rendering in NeRF. Specifically, StyleSDF uses G_0 to predict the distance $d(x) = G_0(w, x)$ between the query point x and the shape surface, where the density function $\sigma(x)$ can be transformed from $d(x)$ for NeRF [27] to render. The incorporation of SDF leads to higher-quality geometry in terms of expressiveness view-consistency and clear definition of the surface. StyleSDF also enjoys the flexible style control for semantic editing as in StyleGAN [21]. Therefore, in this paper we mainly use StyleSDF as the base model for GAN inversion study. Note that our method is not limited to StyleSDF and could be easily extended to other style-based 3D GAN variations.

4. E3DGE

An effective 3D GAN inversion shall be capable of 1) reconstructing plausible 3D shape given single-view input,

2) maintaining high-fidelity texture, and 3) allowing view-consistent semantic edits. To achieve these goals, we propose the E3DGE framework with three novel components: In Sec. 4.1, we leverage 3D GAN to generate pseudo 2D-3D paired samples for 3D supervisions, and train an inversion encoder E_0 to estimate the latent of plausible 3D shapes from a 2D image; In Sec. 4.2, we train a local encoder E_1 to extract pixel-aligned features to enrich texture details for high-fidelity inversion; Finally, Sec. 4.3 introduces a hybrid alignment module for view-consistent semantic editing.

4.1. Self-supervised Inversion Learning

In this section, we propose to mitigate the lack of large-scale high-quality 2D-3D paired datasets by retrofitting pre-trained 3D GANs to provide pseudo samples for training our inversion encoder. We demonstrate the model trained from pseudo samples can rival and even outperform the methods learned from real data on the 3D GAN inversion task. We detail the process as follows.

Global Encoder for 3D GAN Inversion. With the style-based G , we build our encoder E_0 based on pSp [37] for inversion. Given a target image I , E_0 predicts its latent code $\hat{w} = E_0(I)$. Given the corresponding camera pose ξ , the reconstructed image is obtained by $\tilde{I} = G(\hat{w}, \xi)$ to approximate I . In addition, we would like its 3D shape predicted by G_0 to be plausible enough.

Distill 3D GANs as 3D Supervisions. Different compositions of shape and texture could lead to identical 2D-rendered images. 3D supervision is needed to alleviate such shape-texture ambiguity. In the lack of large-scale high-quality 2D-3D paired samples, we formulate GAN Inversion as a *self-training* task, where samples synthesized from itself are leveraged to boost the reconstruction fidelity in both 2D and 3D domains.

As shown in Fig 1, we synthesize paired 3D shape information \mathcal{S} and 2D image I from latent code w and camera pose ξ using G to train E_0 . To extract the 3D shape information \mathcal{S} of each synthetic shape, we first sample a point set $\mathcal{P} = \{\mathcal{P}_O, \mathcal{P}_F\}$ where \mathcal{P}_O and \mathcal{P}_F contain points sampled from the surface and around the surface, respectively. Then, we calculate the geometry descriptor d_i and n_i for each 3D point $x_i \in \mathcal{P}$, and \mathcal{S} is defined as the set of geometry descriptors of all 3D point in \mathcal{P} :

$$\mathcal{S} = \left\{ \left\{ d_i, n_i \right\}_{i=1}^{|\mathcal{P}|} \mid x_i \in \mathcal{P}, d_i = G_0(w, x_i), n_i = \nabla_{x_i} d_i \right\}, \quad (1)$$

where d_i is the distance from x_i to the shape surface and n_i is the surface normal defined by the gradient of the distance w.r.t. x_i . Note our method is not limited to the SDF-based shape representation and can be easily extended to radiance-based methods [7, 8, 33]. Moreover, given different camera poses, we can generate a diverse 2D-3D dataset to help alleviate the shape-texture ambiguity, *i.e.*, for each shape \mathcal{S} ,

various images $\mathbf{I} = G(\mathbf{w}, \xi)$ can be rendered by randomly sampling ξ from a predefined pose distribution p_ξ . Finally, we define $\mathcal{X} = \{\mathcal{S}, \xi, \mathbf{I}\}$ as a training sample for E_0 .

3D GAN-Supervised Training. As shown in Fig. 2 (a), given a training sample \mathcal{X} , the forward process is represented as:

$$\hat{\mathbf{w}} = E_0(\mathbf{I}) \quad (2)$$

$$\{\hat{\mathbf{I}}, \hat{\mathcal{S}}\} = G(\hat{\mathbf{w}}, \xi, \mathcal{P}) \quad (3)$$

where $\hat{\mathbf{w}}$ is the estimated latent code and $\hat{\mathcal{S}} = \{\{\hat{d}_i, \hat{\mathbf{n}}_i\}_{i=1}^{|\mathcal{P}|} \mid \mathbf{x}_i \in \mathcal{P}\}$ is the estimated 3D shape information conditioned on $\hat{\mathbf{w}}$ and \mathcal{P} .

To achieve 3D supervision, we would like the estimated $\hat{\mathcal{S}}$ to approximate the ground truth \mathcal{S} . Specifically, for points over the surface, their distances and normal are both considered while for points around the surface, we only supervise their distance following [3, 35], leading to geometry loss:

$$\mathcal{L}_{geo}^O = \mathbb{E}_{\mathcal{X}} \left[\frac{1}{|\mathcal{P}_O|} \sum_{i=1}^{|\mathcal{P}_O|} \lambda_{g_1} |\hat{d}_i| + \lambda_{g_2} \|\hat{\mathbf{n}}_i - \mathbf{n}_i\|_1 \right] \quad (4)$$

$$\mathcal{L}_{geo}^F = \mathbb{E}_{\mathcal{X}} \left[\frac{1}{|\mathcal{P}_F|} \sum_{i=1}^{|\mathcal{P}_F|} \lambda_{g_3} |\hat{d}_i - d_i| \right] \quad (5)$$

$$\mathcal{L}_{geo} = \mathcal{L}_{geo}^O + \mathcal{L}_{geo}^F, \quad (6)$$

where λ_s are loss weights and $d_i = 0$ for points over the surface. We also impose code reconstruction loss $\mathcal{L}_{code} = \|\hat{\mathbf{w}} - \mathbf{w}\|_2$ to regularize the learning and 2D supervisions \mathcal{L}_{rec} to minimize the reconstruction error between $\hat{\mathbf{I}}$ and \mathbf{I} as in pSp [37]. The overall loss is $\mathcal{L} = \mathcal{L}_{geo} + \mathcal{L}_{code} + \mathcal{L}_{rec}$.

4.2. Local Features for High-Fidelity Inversion

To facilitate introductions in the following sections, we first take a look at the details of StyleSDF. As shown in Fig. 1, G_0 can be further divided into four parts: a 8-layer MLP encoder E_{G_0} , a SDF decoder ϕ_g , a feature decoder ϕ_f and a color decoder ϕ_c . E_{G_0} extracts a global feature $\mathbf{f}_G(\mathbf{x}) = E_{G_0}(\mathbf{x}, \mathbf{w})$. Based on \mathbf{f}_G , ϕ_g and ϕ_f compute SDF $d(\mathbf{x}) = \phi_g(\mathbf{f}_G(\mathbf{x}))$ and the last-layer feature $\mathbf{f}(\mathbf{x}, \mathbf{v}) = \phi_f(\mathbf{f}_G(\mathbf{x}), \mathbf{v})$ of G_0 , respectively. \mathbf{f} could be directly transformed to color $\mathbf{c}(\mathbf{x}, \mathbf{v}) = \phi_c(\mathbf{f}(\mathbf{x}, \mathbf{v}))$ or being volume integrated to \mathbf{F} and sent to G_1 for high resolution synthesis. For simplicity, we will omit \mathbf{v} in the following.

Local Feature for Detailed Textures. The global latent code $\hat{\mathbf{w}}$ is a compact representation of the predicted scene. However, previous works [9, 47] have validated that a low-dimensional latent code discards high-frequency spatial details and fails to reconstruct high-fidelity outputs. This phenomenon becomes more severe when lifting the 2D image to a 3D scene, which contains exponentially more information. Inspired by recent progress in few-shot 3D reconstruction [3, 10, 39, 40, 46, 50, 53], we propose to make up

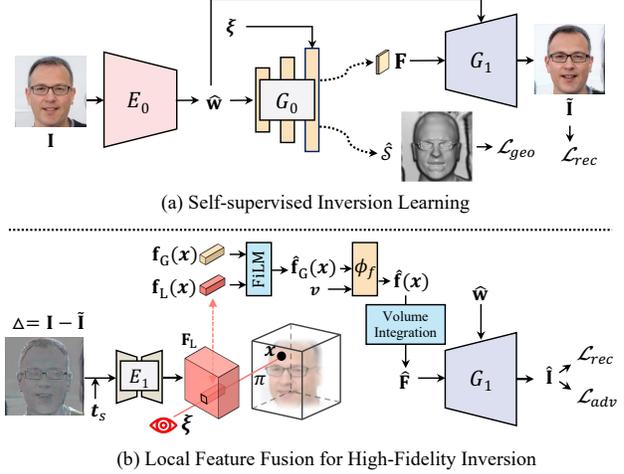


Figure 2. **E3DGE for 3D GAN inversion.** (a) We augment the training of the encoder E_0 with 3D supervision \mathcal{L}_{geo} for plausible 3D shape prediction. (b) We augment the representation capacity of the global latent code $\hat{\mathbf{w}}$ with local point-dependent latent feature \mathbf{f}_L for high-fidelity texture reconstruction.

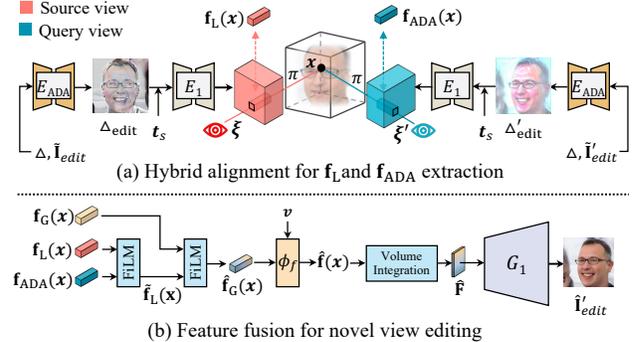


Figure 3. **Hybrid alignment for high-quality editing.** Given code prediction $\hat{\mathbf{w}}$ from encoder E_0 pre-trained in stage-I, we aim to generate high-quality view synthesis over the edited code $\hat{\mathbf{w}}_{edit}$. In (a), the local details Δ along with the target edited image \mathbf{I}_{edit} and depth map $t_s(\hat{\mathbf{w}}, \xi)$ are sent to pre-trained E_{ADA} to predict aligned residual Δ'_{edit} . The original aligned residual Δ along with the 2D auxiliary residual Δ'_{edit} are processed by E_1 to recover latent maps \mathbf{F}_L and \mathbf{F}_{ADA} for later fusion. In (b), the extracted features $\mathbf{f}_L(\mathbf{x})$ and $\mathbf{f}_{ADA}(\mathbf{x})$ are first fused together with a FiLM layer, and the fused result $\hat{\mathbf{f}}_L(\mathbf{x})$ further serve as conditions to modulate the global feature $\mathbf{f}_G(\mathbf{x})$. The final modulated feature $\hat{\mathbf{f}}(\mathbf{x})$ contains complete information, globally and locally. The volume integrated $\hat{\mathbf{F}}$ is sent to G_1 for high-resolution synthesis.

for the lost information by introducing pixel-aligned (local) features. As shown in Fig. 2 (b), rather than conditioning all 3D points with the same latent code $\hat{\mathbf{w}}$, we augment the representation capacity with local latent codes \mathbf{f}_L that is dependent on each point \mathbf{x} . We introduce a local hourglass [28]

encoder E_1 to predict a residual feature map \mathbf{F}_L based on the reconstruction residue $\Delta = \mathbf{I} - \tilde{\mathbf{I}}$,

$$\mathbf{F}_L = E_1(\Delta, t_s(\hat{\mathbf{w}}, \boldsymbol{\xi})), \quad (7)$$

where $t_s(\hat{\mathbf{w}}, \boldsymbol{\xi})$ is the depth map of the scene derived from the SDF to serve as 3D context information. Then, the local latent code of a point \mathbf{x} is its corresponding value in \mathbf{F}_L :

$$\mathbf{f}_L(\mathbf{x}) = \mathbf{F}_L(\pi(\mathbf{x})) \oplus \mathbf{PE}(\mathbf{x}), \quad (8)$$

where π maps the 3D point \mathbf{x} to its corresponding pixel coordinate on 2D feature map \mathbf{F}_L . Since in 3D scenes, points along a ray will be projected to the same coordinate on the 2D plane, to differentiate these points, we additionally concatenate their positional encoding $\mathbf{PE}(\mathbf{x})$ [27] in Eq. (8). In this way, the local feature \mathbf{f}_L only encodes the residual information at the projected position $\pi(\mathbf{x})$ but is also capable of determining where the residual information lies in the 3D scene, as well as inpainting the occluded areas along the ray.

Finally, we fuse the local latent code $\mathbf{f}_L(\mathbf{x})$ with the global latent code $\mathbf{f}_G(\mathbf{x}) = E_{G_0}(\mathbf{x}, \hat{\mathbf{w}})$ to supplement the missing high-frequency details. Specifically, the feature fusion is based on Feature-wise Linear Modulation (FiLM) [36]. As shown in Fig. 2, $\mathbf{f}_L(\mathbf{x})$ is fed into two MLP layers to obtain the scale and bias modulation parameters $\mathbf{f}_L^\gamma(\mathbf{x})$ and $\mathbf{f}_L^\beta(\mathbf{x})$. Then we modulate $\mathbf{f}_G(\mathbf{x})$ with FiLM

$$\hat{\mathbf{f}}_G(\mathbf{x}) = \text{FiLM}(\mathbf{f}_G(\mathbf{x}), \mathbf{f}_L(\mathbf{x})) = \mathbf{f}_L^\gamma(\mathbf{x}) \cdot \mathbf{f}_G(\mathbf{x}) + \mathbf{f}_L^\beta(\mathbf{x}).$$

The fused $\hat{\mathbf{f}}_G(\mathbf{x})$ is volume integrated to $\hat{\mathbf{F}}$ and the final high-fidelity reconstructed image can be obtained as $\hat{\mathbf{I}} = G_1(\hat{\mathbf{F}})$.

Note that through point projection π , the reconstruction with local prior is not limited to the original view, and naturally works for novel views. However, for views with severe occlusions or additional editing, the residual features may not fully align with the scene, leading to a failed feature fusion. We will address this issue in the next subsection with our hybrid feature alignment.

4.3. Hybrid Alignment for High-Quality Editing

Though we achieve high-fidelity reconstruction with the aforementioned designs, there is a trade-off between the *input view* reconstruction quality and *novel view* editing performance. We first analyze the reasons behind and propose a hybrid alignment module to address this issue.

Reconstruction Editing Trade-off. Given an input image \mathbf{I} with paired reconstruction $\tilde{\mathbf{I}}$ and residual map Δ extracted from the input view $\boldsymbol{\xi}$ with the aforementioned method. First, at test time when the input image is edited $\tilde{\mathbf{I}}_{\text{edit}}$ or query view $\boldsymbol{\xi}' \neq \boldsymbol{\xi}$, the residual map no longer aligns and is likely to result in wrong predictions. Second, if we supervise the models to reconstruct the input itself, the learned

features are *regressive* rather than *generative* since all prediction areas are visible in the inputs. With these two challenges, though the model could yield perfect reconstruction at training, it would result in noticeable performance degradation when rendering from novel views at test time.

Hybrid Alignment for High-Quality Editing. To address the first challenge, we propose to infer aligned features with a 2D-3D hybrid alignment. Specifically, given edited latent code $\hat{\mathbf{w}}_{\text{edit}}$, the initial novel-view edited image $\tilde{\mathbf{I}}_{\text{edit}} = G_0(\hat{\mathbf{w}}_{\text{edit}}, \boldsymbol{\xi}')$ is misaligned with Δ . Inspired by HFGI [47], we leverage a 2D alignment module E_{ADA} to address the misalignment. As shown in Fig. 3 (a), we first obtain $\Delta_{\text{edit}} = E_{\text{ADA}}(\Delta, G_0(\hat{\mathbf{w}}_{\text{edit}}, \boldsymbol{\xi}'))$, transform it to residual feature map $\mathbf{F}_L^{\text{edit}}$ via Eq. (7) and retrieve the view-consistent 3D local feature \mathbf{f}_L via Eq. (8). However, to render the high-quality edited image $\tilde{\mathbf{I}}_{\text{edit}}$ from novel view $\boldsymbol{\xi}'$, $\mathbf{F}_L^{\text{edit}}$ might still suffer from occlusion due to large pose variations. To the end, we propose a hybrid alignment to further refine $\mathbf{F}_L^{\text{edit}}$ with 2D aligned feature from E_{ADA} . Specifically, we align a 2D residue $\Delta'_{\text{edit}} = E_{\text{ADA}}(\Delta, \tilde{\mathbf{I}}_{\text{edit}})$ and retrieve its corresponding \mathbf{f}_{ADA} with E_1 , which fills the occlusion in a 2D manner but lacks 3D consistency. To marry the best of both, as shown in in Fig 3 (b), we modulate \mathbf{f}_L with \mathbf{f}_{ADA} ,

$$\tilde{\mathbf{f}}_L(\mathbf{x}) = \text{FiLM}(\mathbf{f}_L(\mathbf{x}), \mathbf{f}_{\text{ADA}}(\mathbf{x})), \quad (9)$$

and further fuse $\tilde{\mathbf{f}}_L$ with $\mathbf{f}_G(\mathbf{x})$ for final prediction,

$$\hat{\mathbf{f}}(\mathbf{x}) = \text{FiLM}(\mathbf{f}_G(\mathbf{x}), \tilde{\mathbf{f}}_L(\mathbf{x})), \quad (10)$$

where $\hat{\mathbf{f}}(\mathbf{x})$ is then integrated to $\hat{\mathbf{F}}$ for rendering the final novel-view edited image $\tilde{\mathbf{I}}_{\text{edit}} = G_1(\hat{\mathbf{F}})$.

Novel View Training for Coherent View Synthesis. To address the second challenge and enforce the model to learn generative features, during training, we sample two views $\boldsymbol{\xi}_1$ and $\boldsymbol{\xi}_2$ for each style code \mathbf{w} , and render the corresponding images $\mathbf{I}^{\boldsymbol{\xi}_1}$ and $\mathbf{I}^{\boldsymbol{\xi}_2}$. Then, we train the models to reconstruct plausible novel views, *i.e.*, $G(E(\mathbf{I}^{\boldsymbol{\xi}_1}), \boldsymbol{\xi}_2) \approx \mathbf{I}^{\boldsymbol{\xi}_2}$ and $G(E(\mathbf{I}^{\boldsymbol{\xi}_2}), \boldsymbol{\xi}_1) \approx \mathbf{I}^{\boldsymbol{\xi}_1}$. This training strategy facilitates a high-quality view synthesis over edited scenes.

5. Experiments

Datasets. We mainly focus on the human face domain and use both 2D and 3D datasets for extensive evaluation. To examine 2D reconstruction quality, we adopt CelebA-HQ [20, 24] dataset for source view reconstruction. To further evaluate novel view reconstruction performance, we synthesize 500 trajectory videos from a pretrained generator as a proxy test set. For attribute editing, we adopt InterfaceGAN [43] and Talk2Edit [19] to search for the editing directions. To evaluate 3D shape reconstruction quality, we use NoW benchmark [41] that provides a rich variety of face images with ground-truth 3D scans. The 3D GANs are

Table 1. Quantitative comparison for inversion quality on faces.

Method	Source View Reconstruction				Novel View Reconstruction			
	MAE ↓	SSIM ↑	LPIPS ↓	Similarity ↑	MAE ↓	SSIM ↑	LPIPS ↓	Similarity ↑
pSp _{StyleSDF}	.150 ± .032	.696 ± .048	.270 ± .059	.498 ± .099	.235 ± .010	.604 ± .011	.358 ± .048	.513 ± .041
e4e _{StyleSDF}	.174 ± .049	.669 ± .049	.226 ± .063	.252 ± .107	.237 ± .014	.597 ± .011	.341 ± .063	.271 ± .060
E3DGE	.097 ± .008	.780 ± .016	.128 ± .017	.883 ± .017	.173 ± .008	.710 ± .010	.154 ± .016	.903 ± .021

pre-trained on FFHQ [21]. Note that our method does not rely on any external 3D data during the training process.

Implementation Details. For all the encoder models, we adopt Adam optimizer with a learning rate of $5e-5$ to train the models on 4 NVIDIA Tesla V100 GPUs, with a resolution of 256^2 , batch size of 24, and 16 samples along a ray for the recommended $200K$ iterations. Following [39], we filter our invisible 3D points when training from a certain view. Code, dataset, and all pre-trained models will be made publicly available. More details are included in the supplementary material.

5.1. Evaluation

5.1.1 Quantitative Evaluation

Since existing baselines are trained on StyleGAN [22] and could be directly applied, for comparison, we implement two canonical encoder-based GAN inversion approaches on StyleSDF [31], *i.e.*, pSp [37] and e4e [44], which stress reconstruction and editing quality respectively.

2D Reconstruction. For 2D evaluation, we report inversion performance for both source view reconstruction and novel view reconstruction in Tab 1. For source view reconstruction, the metrics are calculated on the 2,825 images from CelebA-HQ test set [24]. For novel view reconstruction, the metrics are averaged from 500 videos generated from pre-trained 3D GANs, each with 250 frames covering ellipsoid camera poses trajectory. For each video, we randomly pick one image as source view input and the remaining images as ground truths with labeled poses as query views. In this way, we could extensively evaluate the view synthesis ability under occlusions and varied input viewpoints. Our approach substantially outperforms encoder-based baselines in terms of reconstruction quality in both source view and target view. We include the comparison in the supplementary material and show that our method is considerably faster than optimization-based methods during inference.

3D Reconstruction. We report the 3D face reconstruction performance on NoW benchmark test set in Tab. 2. Our method surpasses purely model-free method [48] and shows competitive performance compared with methods designed for 3D face reconstruction using basic models, *e.g.*, 3DMM [5] and FLAME [25]. Note that as discussed

Table 2. Performance of 3D face reconstruction on NoW [41].

Methods	Prior Type	Median↓	Mean↓	Std
3DMM-CNN [45]	3DMM	1.84	2.33	2.05
PRNet [12]	3DMM	1.50	1.98	1.88
RingNet [41]	FLAME	1.21	1.54	1.31
3DDFA-V2	3DMM	1.23	1.57	1.39
DECA [11]	FLAME	1.09	1.38	1.18
Wu et al. [48]	Model Free	2.64	3.29	2.86
Ours	3D GAN	1.70	2.08	1.67



Figure 4. Qualitative comparisons on face inversions.

in Wu *et al.* [48], NoW benchmark is designed for model-based reconstruction methods and inherently put model-free approaches at a disadvantage. Therefore, our method could serve as a reference for fair quantitative evaluation comparisons of future model-free methods.

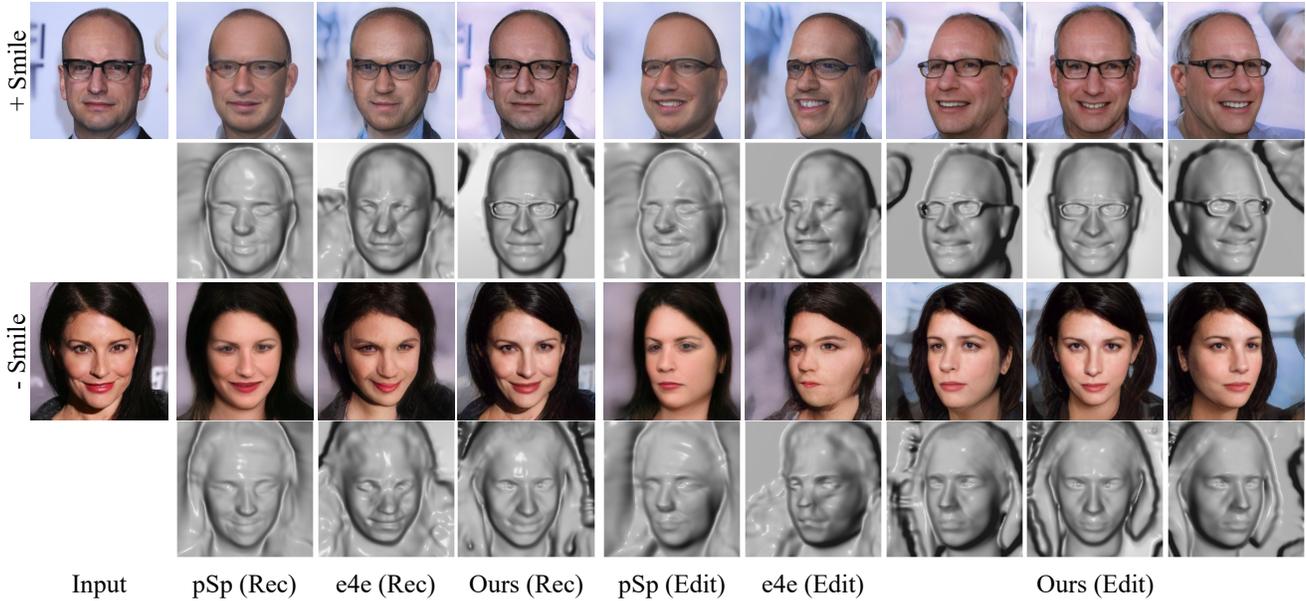


Figure 5. **Qualitative comparisons on face inversion and editing under novel views.**

5.1.2 Qualitative Evaluation

Reconstruction. We show reconstruction performance in Fig. 4. *Geometry-wise*, the baseline models without explicit 3D supervisions tend to generate implausible intermediate shapes, e.g., e4e predictions of rows (3, 6) and pSp predictions of rows (2, 5). Besides, their reconstruction is not close to the “ground truth”, and the reconstructed surface lacks details. Our method successfully regularizes the intermediate 3D shapes and generates plausible results with surface details and a more complete structure. For instance, in rows 4 and 6, our method reconstructs 3D eyeglasses in which the baselines fail. Corresponding metrics in Tab. 4 also validate the usefulness of the direct geometry supervisions and loss designs. *Texture-wise*, existing methods generate distorted results and suffer artifacts and identity change. In contrast, with pixel-aligned features incorporated, our method is more robust with high-fidelity results. In particular, our method captures more details and preserves the identity of different input viewpoints. For example, in row 1 – 3, our method accurately reconstructs the hair, and in row 5, the beard.

Editing. We include the editing results in Fig. 5 and choose the “Smile” attribute for editing. Beyond plausible shape reconstruction with high-fidelity texture inversion, in-view synthesis over edited results, our method consistently generates high-quality edited renderings in terms of view consistency, details conservation, and identity preservation. Compared with our method, the baselines either fail to render intact identity (column 5) or generate visually plausible shapes (column 6).

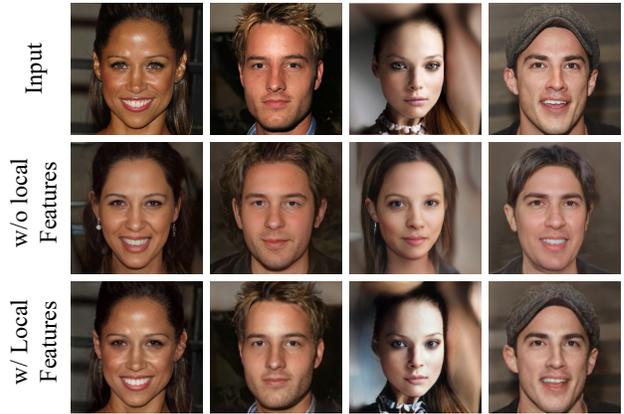


Figure 6. **Ablation of Local Features.** Our method with pixel-aligned features shows photorealistic reconstructions..

5.2. Ablation Study

Effect of 3D GAN as Supervisions. We quantitatively validate the effects of 3D supervision in NoW Challenge validation set and report the corresponding metrics in Tab. 4. For the results of fully synthetic dataset training (row 1), compared with the baseline method with a similar network (pSp), fully synthetic data training shows worse reconstruction metrics. We attribute this phenomenon to the domain gap between synthesized images and real images. However, our method shows surprisingly better performance over identity preservation in novel views (0.77 compared with 0.513 of pSp_{StyleSDF} and 0.271 of e4e_{StyleSDF}, which we

Table 3. **Ablations of Local Features and Hybrid Fusion.** Our local-global model design with hybrid alignment achieves the balance of high-quality reconstruction and view synthesis.

Ablation Settings	Source View Reconstruction				Novel View Reconstruction			
	MAE ↓	SSIM ↑	LPIPS ↓	Similarity ↑	MAE ↓	SSIM ↑	LPIPS ↓	Similarity ↑
Synthetic Training	.245 ± .024	.634 ± .019	.333 ± .029	.369 ± .056	.241 ± .011	.594 ± .008	.366 ± .059	.770 ± .026
+Local Features	.074 ± .007	.811 ± .015	.075 ± .010	.953 ± .006	.282 ± .103	.571 ± 0.056	.511 ± 0.031	.608 ± .123
+3D Alignment	.102 ± .009	.772 ± .015	.119 ± .016	.818 ± .029	.133 ± .011	.709 ± .022	.130 ± .021	.901 ± .011
+2D Alignment	.098 ± .005	.774 ± .038	.140 ± .040	.900 ± .032	.178 ± .007	.656 ± .009	.178 ± .012	.904 ± .018
Hybrid Alignment	.097 ± .008	.780 ± .016	.128 ± .017	.883 ± .017	.131 ± .008	.710 ± .010	.154 ± .016	.903 ± .021

Table 4. Effect of 3D Supervisions.

Settings	Median ↓	Mean ↓	Std
pSp _{StyleSDF}	1.97	2.43	2.05
e4e _{StyleSDF}	2.83	3.40	2.67
+ \mathcal{L}_{geo}^O	1.75	2.11	1.72
+ \mathcal{L}_{geo}^F	1.71	2.09	1.70
+ \mathcal{L}_{code}	1.66	2.06	1.69

attribute to the well-aligned pose of synthetic corpus leads to less distortion in view synthesis.

Effect of Local Features. As discussed before, the local features preserve the missing image details to facilitate high-fidelity reconstruction. To validate the effectiveness of local features in texture reconstructions, we show the inversion results in Fig. 6. With the proposed local-global fusion pipeline, our model captures more details and guarantees photorealistic reconstruction. Quantitative results in Tab. 3 also validate the effectiveness of local features in high-quality inversion. The results on the video trajectories also show that without delicate design, *e.g.* novel-view training, local features would fully collapse over novel view synthesis.

Effect of Hybrid Alignment. We show the view synthesis achieved by different alignment methods in Fig. 7. To quantitatively analyze the effect of hybrid alignment, in Tab. 3 we evaluate the model performance of 3D alignment and 2D alignment individually. For both ablations, novel-view training is enabled. As shown here, the 3D alignment model shows better view consistency in video prediction measured by reconstruction metrics, and the 2D alignment model shows better identity preservation. The hybrid alignment model marries the best of both and also enables semantic editing and yields better reconstruction performance on the video predictions.

6. Conclusion and Discussions

We propose a novel 3D GAN inversion framework E3DGE for 3D face reconstruction and editing. We marry the benefits of both self-supervised global prior and pixel-

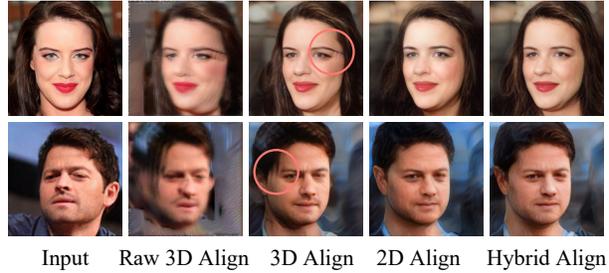


Figure 7. **Ablation of Hybrid Alignment.** From left to right, we show the novel view synthesis of raw 3D-aligned features w/o novel-view training, synthesis achieved using 2D-aligned features, and the final hybrid features. 3D-aligned features are view-consistent but suffer from occlusions (circled), while 2D features are visually plausible but lack some details (*e.g.*, hair color). Our hybrid fused results share the best of both.

aligned local prior for high-quality shape and texture reconstruction. A hybrid alignment that bridges the best of 2D and 3D features is further proposed for view-consistent editing. Benefiting from the overall system design, the proposed method has advantages in terms of both high fidelity and editability. As a pioneer attempt in this direction, we believe this work opens a new line of research direction and will inspire future works on 3D GAN inversion, few-shot 3D reconstruction and 3D-aware learning from 2D images.

Limitations and Future Work. The proposed method suffers data bias introduced by the synthetic data. As the synthetic data lacks complex details and pose variations compared with real-world data, our method trained with it tends to generate simple background and fail on extreme poses. Special attentions should be paid to data bias to avoid social impact to under represented minorities. A future direction is to leverage real data for semi-supervised training. Another future direction is to leverage the hyper-network for efficient local feature incorporation to alleviate the extra computational cost of the 2D alignment module. Finally, we would explore the potentials of our framework on other 3D GANs and shapes beyond human face and other editing methods uniquely designed for 3D GANs.

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*, 2019. 1, 2
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2StyleGAN++: How to edit the embedded images? In *CVPR*, 2020. 1, 2
- [3] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3D reconstruction of humans wearing clothing. *CVPR*, 2022. 4
- [4] Victor Besnier, Himalaya Jain, Andrei Bursuc, Matthieu Cord, and Patrick Pérez. This Dataset Does Not Exist: Training Models from Generated Images. *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020. 2
- [5] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, 1999. 6
- [6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*. OpenReview.net, 2019. 2
- [7] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and G. Wetzstein. Pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis. In *CVPR*, 2021. 2, 3
- [8] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022. 2, 3
- [9] Kelvin C.K. Chan, Xiangyu Xu, Xintao Wang, Jinwei Gu, and Chen Change Loy. GLEAN: Generative latent bank for large-factor image super-resolution and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2, 4
- [10] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo Radiance Fields (SRF): Learning View Synthesis for Sparse Views of Novel Scenes. *CVPR*, 2021. 4
- [11] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. In *SIGGRAPH*, volume 40, 2021. 6
- [12] Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. Joint 3d face reconstruction and dense alignment with position map regression network. In *ECCV*, 2018. 6
- [13] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [14] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleNeRF: A style-based 3D-aware generator for high-resolution image synthesis. In *ICLR*, 2021. 2, 3
- [15] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Escaping plato’s cave: 3D shape from adversarial rendering. In *ICCV*, 2019. 2
- [16] Fangzhou Hong, Zhaoxi Chen, Yushi Lan, Liang Pan, and Ziwei Liu. Eva3d: Compositional 3d human generation from 2d image collections. *arXiv preprint arXiv:2210.04888*, 2022. 2
- [17] Ali Jahanian, Lucy Chai, and Phillip Isola. On the” steerability” of generative adversarial networks. *The International Conference on Learning Representations (ICLR)*, 2020. 2
- [18] Ali Jahanian, Xavier Puig, Yonglong Tian, and Phillip Isola. Generative models as a data source for multiview representation learning. *ICLR*, 2022. 2
- [19] Yuming Jiang, Ziqi Huang, Xingang Pan, Chen Change Loy, and Ziwei Liu. Talk-to-Edit: Fine-grained facial editing via dialog. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 5
- [20] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *ICLR*, 2018. 5
- [21] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2, 3, 6
- [22] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 2, 6
- [23] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 3
- [24] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. MaskGAN: Towards diverse and interactive facial image manipulation. In *CVPR*, 2020. 5, 6
- [25] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *TOG*, 36(6), 2017. 6
- [26] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *CVPR*, June 2019. 2
- [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 *ECCV*. Springer, 2020. 2, 3, 5
- [28] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *ECCV*, 2016. 4
- [29] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yongliang Yang. HoloGAN: Unsupervised Learning of 3D Representations From Natural Images. In *ICCV*, 2019. 2
- [30] Michael Niemeyer and Andreas Geiger. GIRAFFE: Representing scenes as compositional generative neural feature fields. In *CVPR*, 2021. 2, 3
- [31] 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 *CVPR*, 2021. 2, 3, 6
- [32] Xingang Pan, Bo Dai, Ziwei Liu, Chen Change Loy, and Ping Luo. Do 2D GANs know 3D shape? Unsupervised 3D Shape Reconstruction from 2D Image GANs. In *ICLR*, 2021. 2

- [33] Xingang Pan, Xudong Xu, Chen Change Loy, Christian Theobalt, and Bo Dai. A Shading-Guided Generative Implicit Model for Shape-Accurate 3D-Aware Image Synthesis. In *NIPS*, 2021. 3
- [34] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting Deep Generative Prior for Versatile Image Restoration and Manipulation. *TPAMI*, 44:7474–7489, 2022. 2
- [35] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*. IEEE, 2019. 2, 4
- [36] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. FiLM: Visual reasoning with a general conditioning layer. In *AAAI*, 2018. 5
- [37] 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 *CVPR*, 2021. 1, 2, 3, 4, 6
- [38] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Trans. Graph.*, 2021. 1
- [39] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, October 2019. 4, 6
- [40] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *CVPR*, June 2020. 4
- [41] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael Black. Learning to regress 3D face shape and expression from an image without 3D supervision. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 5, 6
- [42] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative radiance fields for 3D-aware image synthesis. In *NIPS*, 2020. 2
- [43] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. InterFaceGAN: Interpreting the disentangled face representation learned by GANs. *PAMI*, PP, 2020. 5
- [44] 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. 1, 2, 6
- [45] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gerard Medioni. Regressing robust and discriminative 3d morphable models with a very deep neural network. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [46] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P. Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas A. Funkhouser. IBRNet: Learning Multi-View Image-Based Rendering. In *CVPR*, pages 4688–4697, 2021. 4
- [47] Tengfei Wang, Yong Zhang, Yanbo Fan, Jue Wang, and Qifeng Chen. High-Fidelity GAN inversion for image attribute editing. In *CVPR*, 2022. 1, 2, 4, 5
- [48] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised Learning of Probably Symmetric Deformable 3D Objects from Images in the Wild. In *CVPR*, 2020. 6
- [49] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. GAN Inversion: A Survey. *TPAMI*, 2022. 1
- [50] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J. Black. ICON: Implicit Clothed humans Obtained from Normals. In *CVPR*, June 2022. 4
- [51] Yinghao Xu, Sida Peng, Ceyuan Yang, Yujun Shen, and Bolei Zhou. 3d-aware image synthesis via learning structural and textural representations. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18409–18418, 2021. 2
- [52] Shuai Yang, Liming Jiang, Ziwei Liu, , and Chen Change Loy. VToonify: Controllable high-resolution portrait video style transfer. *ACM Transactions on Graphics (TOG)*, 41(6):1–15, 2022. 2
- [53] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. PixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. 4
- [54] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. DatasetGAN: Efficient labeled data factory with minimal human effort. In *CVPR*, 2021. 2
- [55] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain GAN Inversion for Real Image Editing. In *ECCV*, 2020. 2

Self-Supervised Geometry-Aware Encoder for Style-Based 3D GAN Inversion

Supplementary Material

A. Background

Since recent 3D-aware image generative models are all based on neural implicit representations, especially NeRF [12], here we briefly introduce the NeRF-based 3D representation and more StyleSDF details for clarification.

NeRF-based 3D Representation. NeRF [12] proposed an implicit 3D representation for novel view synthesis. Specifically, NeRF defines a scene as $\{c, \sigma\} = F_{\Phi}(\mathbf{x}, \mathbf{v})$, where \mathbf{x} is the query point, \mathbf{v} is the viewing direction from camera origin to \mathbf{x} , c is the emitted radiance (RGB value), σ is the volume density. To query the RGB value $C(\mathbf{r})$ of a point on a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{v}$ shoot from the 3D coordinate origin \mathbf{o} , we have the volume rendering formulation,

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))c(\mathbf{r}(t), \mathbf{v})dt, \quad (1)$$

where $T(t) = \exp(-\int_{t_n}^t \sigma(\mathbf{r}(s))ds)$ is the accumulated transmittance along the ray \mathbf{r} from t_n to t . t_n and t_f denote the near and far bounds.

More StyleSDF Details. In hybrid 3D generation [3, 7, 13], the intermediate feature map is calculated by replacing the color c with feature \mathbf{f} from ϕ_f , namely $\mathbf{F}(\mathbf{r}) = \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))\mathbf{f}(\mathbf{r}(t), \mathbf{v})dt$. In StyleSDF, the Sigmoid activation function σ is replaced by $\sigma(\mathbf{x}) = K_{\alpha}(d(\mathbf{x})) = \text{Sigmoid}(-d(\mathbf{x})/\alpha)/\alpha$, where α is a learned parameter that controls the tightness of the density around the surface boundary.

Notation Table. For clarity, we include the notations used in the proposed method in Tab. 1.

B. Implementation Details

B.1. More Methods Details

Surface Point Sampling in Self-supervised Inversion Learning. In Sec. 4.1 of the main paper, to extract the 3D shape information \mathcal{S} of each synthetic shape, we first sample a point set $\mathcal{P} = \{\mathcal{P}_{\mathcal{O}}, \mathcal{P}_{\mathcal{F}}\}$ where $\mathcal{P}_{\mathcal{O}}$ and $\mathcal{P}_{\mathcal{F}}$ contain points sampled from the surface and around the surface, respectively. To get points over the surface $\mathcal{P}_{\mathcal{O}}$ for training, for efficiency, we directly reuse the intermediate results to render $\mathbf{I}_{\mathcal{O}}$ to calculate the surface. Specially, to sample point

set \mathcal{O} we replace the color c as the coordinates \mathbf{x} of points along a ray in Eq. (1) and approximate the 3D coordinates of surface, namely $\mathbf{t}_s(\mathbf{w}, \xi) = \int_{t_n}^{t_f} T(t, \mathbf{w})\sigma(\mathbf{r}(t), \mathbf{w})t dt$. In this way, we get $B \times H \times W$ surface points for training in each iteration, where B stands for batch size and $H \times W$ stands for the resolution to render 3D consistent images, e.g., 64×64 . To sample point set \mathcal{F} , we add Gaussian offset to each of the calculated surface points \mathcal{O} . Specifically, we adopt Gaussian distribution $\mathcal{N}(0, (r/4)^2)$ where r is the radius of the scene. In this way, points falling within 4 standard deviations would cover 95.44% of the whole 3D space. Following PIFu [19], we also uniformly sample $0.5 \times B \times H \times W$ points within the whole 3D space defined. The overall quantity of the point set surface is $|\mathcal{F}| = 1.5 \times B \times H \times W$. We find this sampling strategy avoids overfitting and yields better performance.

Training Details of High-Fidelity Inversion With Local Features. In Sec. 4.2 of the main paper, we train a local encoder E_1 to extract pixel-aligned features to enrich texture details for high-fidelity inversion. The network architecture of E_1 is identical to that of PIFu [19], which is a stacked hourglass network with residual connections. The input residual map resolution is 256×256 , and the output 64×64 resolution feature map. $\mathbf{f}_L \in \mathcal{R}^{256}$ is bilinearly interpolated from feature map \mathbf{F}_L at the projected position $\pi(\mathbf{x})$. As shown in Fig. 1, we implement the FiLM layer [14] with two MLP residual blocks [28], which outputs α and β for modulation, respectively. We use the identical learning rate and optimizer to train E_1 .

Novel-View Training Details. For novel-view training for coherent view synthesis in Sec. 4.3 of the main paper, in each training iteration with batch size n , rather than sampling n different latent codes $\{\mathbf{z}_i\}_{i=1}^n$, we halve the number of identical latent codes $\{\mathbf{z}_i\}_{i=1}^{n/2}$ while double the rendered images for each latent code $\{\mathbf{I}_i^{\xi_1}, \mathbf{I}_i^{\xi_2}\}_{i=1}^{n/2}$ where n is even. Thus, we train the models to reconstruct plausible *novel views*, i.e., $G(E(\mathbf{I}_i^{\xi_1}), \xi_2) \approx \mathbf{I}_i^{\xi_2}$ and $G(E(\mathbf{I}_i^{\xi_2}), \xi_1) \approx \mathbf{I}_i^{\xi_1}$. Since the paired-sampled images could serve as both inputs and ground truths, the effective batch size and training cost maintains the same. To train 2D alignment model E_{ADA} , we further regularize the predicted residual map $\hat{\Delta}^{\xi_1} \approx \mathbf{I}^{\xi_1} - \mathbf{I}_0^{\xi_1}$ with \mathcal{L}_1 loss, where $\mathbf{I}_0^{\xi_1}$ is correspond-

ing renderer output low-resolution image and $\lambda_1 = 0.1$. Note that we finetune pre-trained E_{ADA} from HFGI [24] with novel-view training and no edited images are involved in the training time.

Curriculum Pose Sampling. At the beginning of the training of the hybrid alignment in Sec. 4.3 of the main paper, large view changes will make the prediction of residual features and the inpainting of occlusion regions extremely difficult. As a result, our model is prone to blurry results. We attribute the reason to the ill-posed nature of rendering novel views given partial observations since the inpainted image is not unique. To facilitate novel-view training, we design a curriculum learning strategy [6] based on *pose sampling difficulty*. Implementation wise, given the camera pose distribution $\xi \sim p_\xi$ with mean μ and standard variance σ , we fix the μ and scale the σ with a weight α which is initially set to 0 and gradually increases to 1 as the training goes. Intuitively, when $\alpha = 0$ the source view ξ is identical to the query view ξ' , the training degrades to a regression task where the model shall reconstruct all the texture details to minimize the loss. As the variance $\alpha \cdot \sigma$ increases, the training becomes a conditional generation task to inpaint plausible and photo-realistic areas.

B.2. More Experiments Details

Training Details. In this work, we directly use the officially released pre-trained GAN models from StyleSDF. In self-supervised shape inversion learning (Sec. 4.1), due to GPU memory restriction, we sample 4 shapes per GPU each iteration for training. After E_0 converged, we fix the network weights and only train the E_1 for high-fidelity inversion. We train each stage for 50,000 iterations, which costs 2 days on 4 Tesla V100 GPUs.

Network Architecture Details. For E_0 , a modified version of the pSp encoder [17] is deployed here for a fair comparison with existing work. Since G_0 and G_1 of StyleSDF have 9 and 10 latent codes, respectively, we introduce $9 + 10$ extra prediction heads to the pSp for the latent code prediction. We observe that early layers of G_0 control the geometry of generated samples, and later G_0 layers as well as decoder generator G_1 control the texture and high-frequency details. Thus, we adopt the early pSp feature map of resolution 32×32 to predict latent code of G_0 for geometry control, and pSp feature map of resolution 64×64 to predict latent code of G_0 for texture control. We use the highest resolution feature map of pSp with resolution 128×128 to predict the latent code for G_1 . We show our FiLM layer implementation in Fig. 1, where the input features are modulated by the input conditions with predicted γ , and β . The MLP is implemented with the MLP residual block [28].

Editing. For attribute editing, following previous works, we adopt vector-arithmetic [16] based editing. Specifically, a searched latent code vector paired with a certain attribute

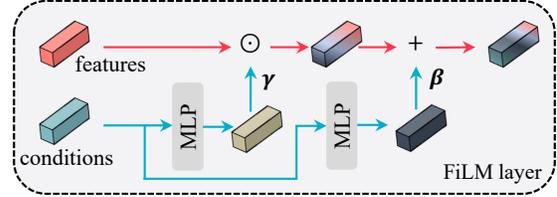


Figure 1. FiLM Layer Architecture.

is weighted and added to the predicted code \hat{w} . To search for the meaningful editing directions on the 3D GAN used, we first sample 10,000 images with paired latent codes from StyleSDF, and then apply the face attribute predictor from Talk-to-Edit [8] to predict the corresponding attributes score. Based on the prediction, we apply SVM classifier from InterfaceGAN [22] to search for the decision boundary. As in previous works [17, 23], we search for the editing latent code in the \mathcal{W} space.

3D Face Reconstruction Evaluation Details. We evaluate the reconstructed 3D meshes and compare them with the performance of several model-based reconstruction methods on NoW benchmark [20]. NoW benchmark [20], provides a test set of 1,702 images of 80 subjects and a ground-truth 3D scan per subject. These images are captured with a higher variety in facial expression, occlusion, and lighting and shall validate the generality of single-view reconstruction methods under real-world conditions.

To extract meshes for evaluation, we detect faces and crop the images using RetinaFace [21] implemented by [25] and obtain 3D mesh reconstructions from the depth maps predicted by our method trained on FFHQ pre-trained generator. We then use the evaluation protocol provided by the benchmark, which aligns the predicted meshes with the ground-truth meshes with a rigid transformation based on seven pre-defined keypoints and computes the scan-to-mesh distances. We obtain keypoints on our predicted meshes by applying a facial keypoint detector [26] on the reconstructed canonical images. Following Unsup3D [27], the average keypoints are used when the keypoint detector fails.

Video Trajectory Evaluation Details. We sample 500 trajectory videos with pre-trained FFHQ StyleSDF generator with an ellipsoid trajectory of size 250 from official StyleSDF code, making a dataset of size 12,5000. The evaluation code and dataset will be released.

B.3. Losses

Reconstruction Loss. We briefly introduce the supervisions we adopt in image reconstructions in both training stages. First, we utilize the pixel-wise \mathcal{L}_2 loss,

$$\mathcal{L}_2(\mathbf{I}) = \|\mathbf{I} - \hat{\mathbf{I}}\|_2. \quad (2)$$

In addition, to learn perceptual similarities, we use the LPIPS [29] loss, which has been shown to better preserve

image quality compared to the more standard perceptual loss:

$$\mathcal{L}_{\text{LPIPS}}(\mathbf{I}) = \|F(\mathbf{I}) - F(\hat{\mathbf{I}})\|_2, \quad (3)$$

where $F(\cdot)$ denotes the perceptual feature extractor.

Finally, a common challenge when handling the specific task of encoding facial images is the preservation of the input identity. To tackle this, we incorporate a dedicated recognition loss measuring the cosine similarity between the output image and its source,

$$\mathcal{L}_{\text{Similarity}}(\mathbf{I}) = 1 - \langle R(\mathbf{I}), R(E_g(\mathbf{I})) \rangle, \quad (4)$$

where R is the pretrained ArcFace [4] network.

In summary, the total loss function is defined as

$$\mathcal{L}_{\text{rec}}(\mathbf{I}) = \lambda_1 \mathcal{L}_2(\mathbf{I}) + \lambda_2 \mathcal{L}_{\text{LPIPS}}(\mathbf{I}) + \lambda_3 \mathcal{L}_{\text{Similarity}}(\mathbf{I}),$$

where we set $\lambda_1 = 1$, $\lambda_2 = 0.8$, $\lambda_3 = 0.1$ as the defined loss weights. In E_0 training, we supervise images $\hat{\mathbf{I}}_0, \hat{\mathbf{I}}_1$ of both resolutions. In E_1 training, we only supervise the reconstruction of high-resolution images since the network weights to render $\hat{\mathbf{I}}_0$ is fixed. Here, we also impose the non-saturating adversarial loss with R1 regularization [11] to improve the naturalness of reconstructed images, which is defined as:

$$\mathcal{L}_{\text{adv}} = -\mathbb{E}[\log(D(\hat{\mathbf{I}}))], \quad (5)$$

$$\mathcal{L}_D = \mathbb{E}[\log(D(\hat{\mathbf{I}}))] + \mathbb{E}[\log(1 - D(\mathbf{I}))], \quad (6)$$

$$\mathcal{L}_{R1} = \lambda \|\nabla D(\hat{\mathbf{I}}; \theta_D)\|_2, \quad (7)$$

where D is initialized with the pre-trained discriminator paired with the generator and θ_D is the corresponding parameters to optimize. In summary, the overall loss is the weighted summation of the loss functions described above:

$$\mathcal{L} = \mathcal{L}_{\text{geo}} + \mathcal{L}_{\text{rec}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}} + \lambda_D \mathcal{L}_D + \lambda_{R1} \mathcal{L}_{R1}, \quad (8)$$

where we set $\lambda_D = \lambda_{\text{adv}} = 0.01$ and $\lambda_{R1} = 10$ in the experiments.

C. More Results

Comparisons with Optimization-based Methods. We include the comparisons with two canonical optimization-based methods here, namely SG2 [1, 9] which is initially proposed in StyleGAN [9] paper to project input image to the \mathcal{W} space of the paired generator, and PTI [18] which further finetune the generator weights to achieve high-fidelity inversion. We implement SG2 and PTI following the official implementations and tune the corresponding parameters for StyleSDF generator. For SG2, we optimize 450 steps with learning rate $5e - 3$, and for the pivotal tuning stage, we optimize 100 steps with learning rate $5e - 5$. We will release all inversion-related code upon acceptance.

We show the qualitative comparison in Fig. 2. As can be seen, SG2 could not reconstruct high-fidelity texture details but maintains a plausible intermediate shape inversion, due to the strong regularization of \mathcal{W} space. Though PTI could achieve photorealistic reconstruction, it still could not alleviate the shape-texture ambiguity, leaving the inverted shape distorted.

We also include the quantitative comparisons in Tab. 2. Specifically, for 2D inversion metrics, we inverse each image in the test set (2,780 CelebA-HQ images) with SG2 and PTI and calculate the reconstruction metrics as well as the inference time. For 3D inversion metrics, we adopt the NoW challenge validation set and reconstruct the corresponding depth mesh for 352 identities. As can be seen, SG2 cannot achieve high-fidelity reconstruction, and PTI could yield high-quality reconstruction at the cost of inference time and shape quality. Our proposed method achieves a balance of both and holds the merit of speedy inference, with only 0.19 seconds needed to render an image from a novel view.

More Comparisons with Encoder-based Methods. Here, we include more comparisons with encoder-based methods in Fig. 3. Our method achieves consistently better performance compared to the baselines in terms of reconstruction fidelity and editing visual quality.

More Editing Results. We show more editing results on changing 4 semantic attributes of our proposed method, namely smile (Fig. 4), hair/beard (Fig. 5), age (Fig. 6) and bangs (Fig. 7). Our method shows promising performance with shape-texture consistent editing. Note that since StyleSDF is still built on an MLP-based generator [2] and InterfaceGAN [22] is also not designed for 3D GANs, the editing performance is hindered to some extent and cannot achieve comparable performance compared with 2D StyleGAN. However, we believe this limitation could be alleviated in the future by adopting better-designed 3D GAN architecture, e.g., tri-plane [3] and vision transformer [5]. Our results unleash the potential of this field and show that 3D consistency and high-fidelity reconstruction with high-quality editing are also achievable in recently developed 3D GAN. We hope our method could inspire later work in this field.

More Toonify Results. We show 3D toonify-stylized results over real-world faces using our proposed method in Fig. 8. Following [15], we finetune the pre-trained generator G for 400 iterations with 317 cartoon face images and use our pre-trained encoder E for inference. Visually inspected, the toonified results holds the cartoon style and also preserve identity of the input image, which demonstrates the potential of applying our method over downstream tasks.

Table 1. Notations used in the proposed method.

Notation	Meaning
$\hat{*}$	Final predictions
$*^*$	Intermediate results
$*'$	Abbreviation of target view camera pose
G	Generator
G_0	Renderer Generator
G_1	SR Generator
D	Discriminator
E	Encoder
E_0	Encoder to predict global latent code
E_1	Hourglass encoder to predict pixel-aligned local features.
E_{ADA}	ADA (Adaptive Distortion Alignment) module
\mathcal{W}	W space for style-based GAN
\mathbf{w}	Latent code sampled from W space
\mathbf{I}	Input image
\mathbf{I}_0	Rendered image from renderer generator
\mathbf{I}_{edit}	Edited image
$\hat{\mathbf{w}}$	Predicted latent code from E_0
λ	Loss weights
\mathbf{x}	3D point
\mathcal{P}	Point set
\mathcal{P}_O	Point set sampled from object surface
\mathcal{P}_F	Point set sampled near the surface or uniformly in the defined 3D space.
d	Signed distance function
\mathbf{n}	Normal for a point
ϕ_g	MLP to predict geometry
ϕ_f	MLP to predict view-dependent feature
ϕ_c	MLP to predict color
\mathbf{v}	View direction
\mathcal{X}	A synthetic data sample for training
ξ	Source view camera pose
ξ'	Target view camera pose
Δ	Residual of predicted image and input image
Δ_{edit}	Residual paired with an edited image
Δ'_{edit}	Residual paired with an edited image rendered from target camera pose.
$\pi(\mathbf{x})$	Projection of 3D point \mathbf{x} to source view
\oplus	Concatenation
PE	Positional Encoding
β, γ	Modulation signals for FiLM
$\mathbf{t}_s(\mathbf{w}, \xi)$	Depth map for code \mathbf{w} rendered from pose ξ
F	Feature map
F_L	Local feature map output from E_1
$\hat{\mathbf{F}}$	Modulated feature map for final prediction
F_{ADA}	Local feature map output from E_1 with E_{ADA} aligned residual
f_G	Global feature output from the generator.
f_L	Local feature interpolated from F_L
f_{ADA}	Aligned feature interpolated from F_L
$\hat{\mathbf{f}}_L$	Predicted local feature for final prediction

Table 2. Quantitative comparisons with optimization-based methods on faces.

Method	CelebA-HQ [10]				NoW Challenge [20] Validation Set			Inference Time
	MAE ↓	SSIM ↑	LPIPS ↓	Similarity ↑	Median ↓	Mean ↓	Std	Second (s) ↓
SG2 [9]	.202 ± .063	.650 ± .054	.167 ± .046	.219 ± .106	1.89	2.23	1.82	235s
PTI [18]	.062 ± .012	.796 ± .017	.027 ± .005	.892 ± .009	2.86	3.54	3.01	265s
E3DGE	.097 ± .008	.780 ± .016	.128 ± .017	.883 ± .017	1.66	2.06	1.69	0.19s (Texture) / 0.81s (Shape)

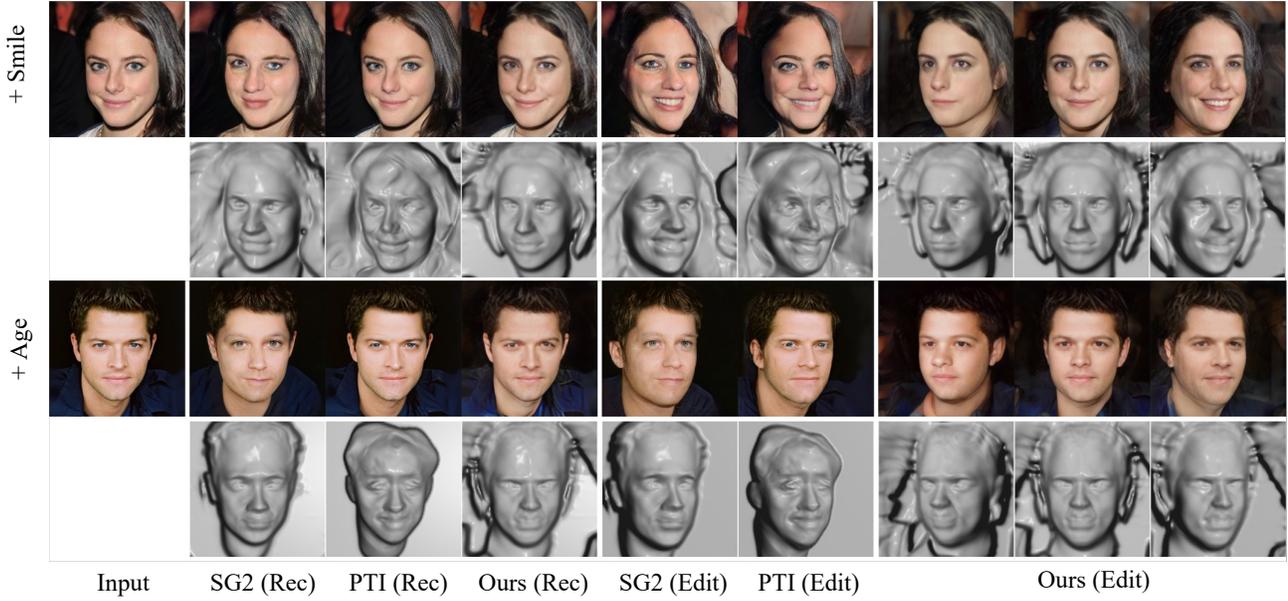


Figure 2. Visual comparisons on optimization-based methods. 'Rec' and 'Edit' represent reconstruction and editing, respectively.

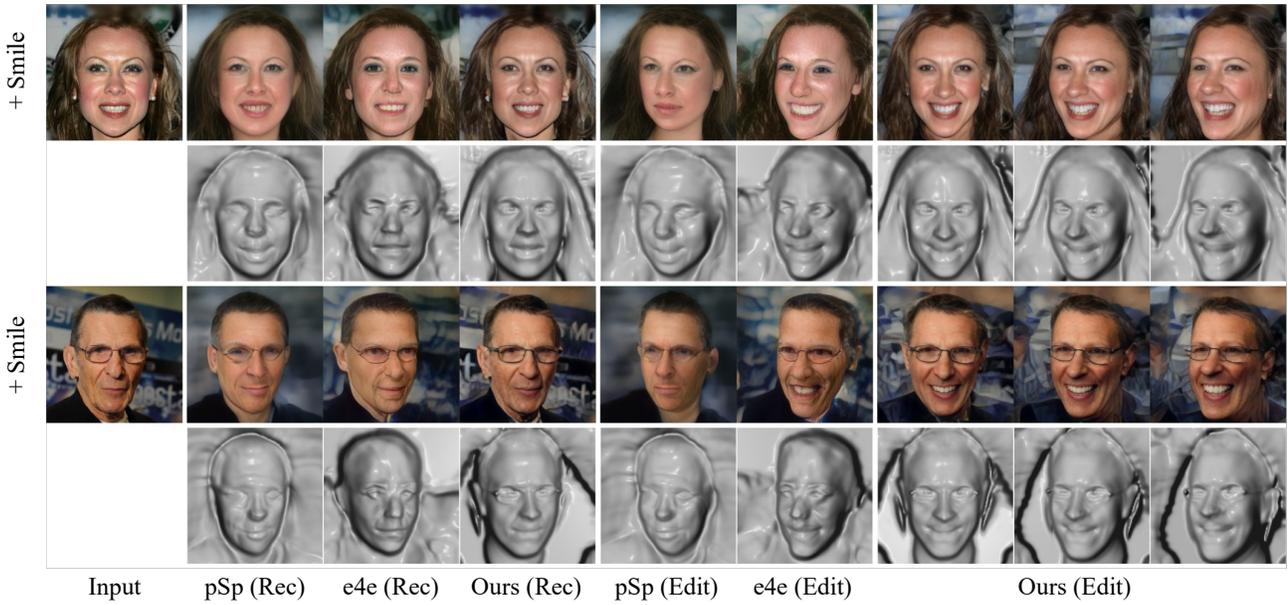
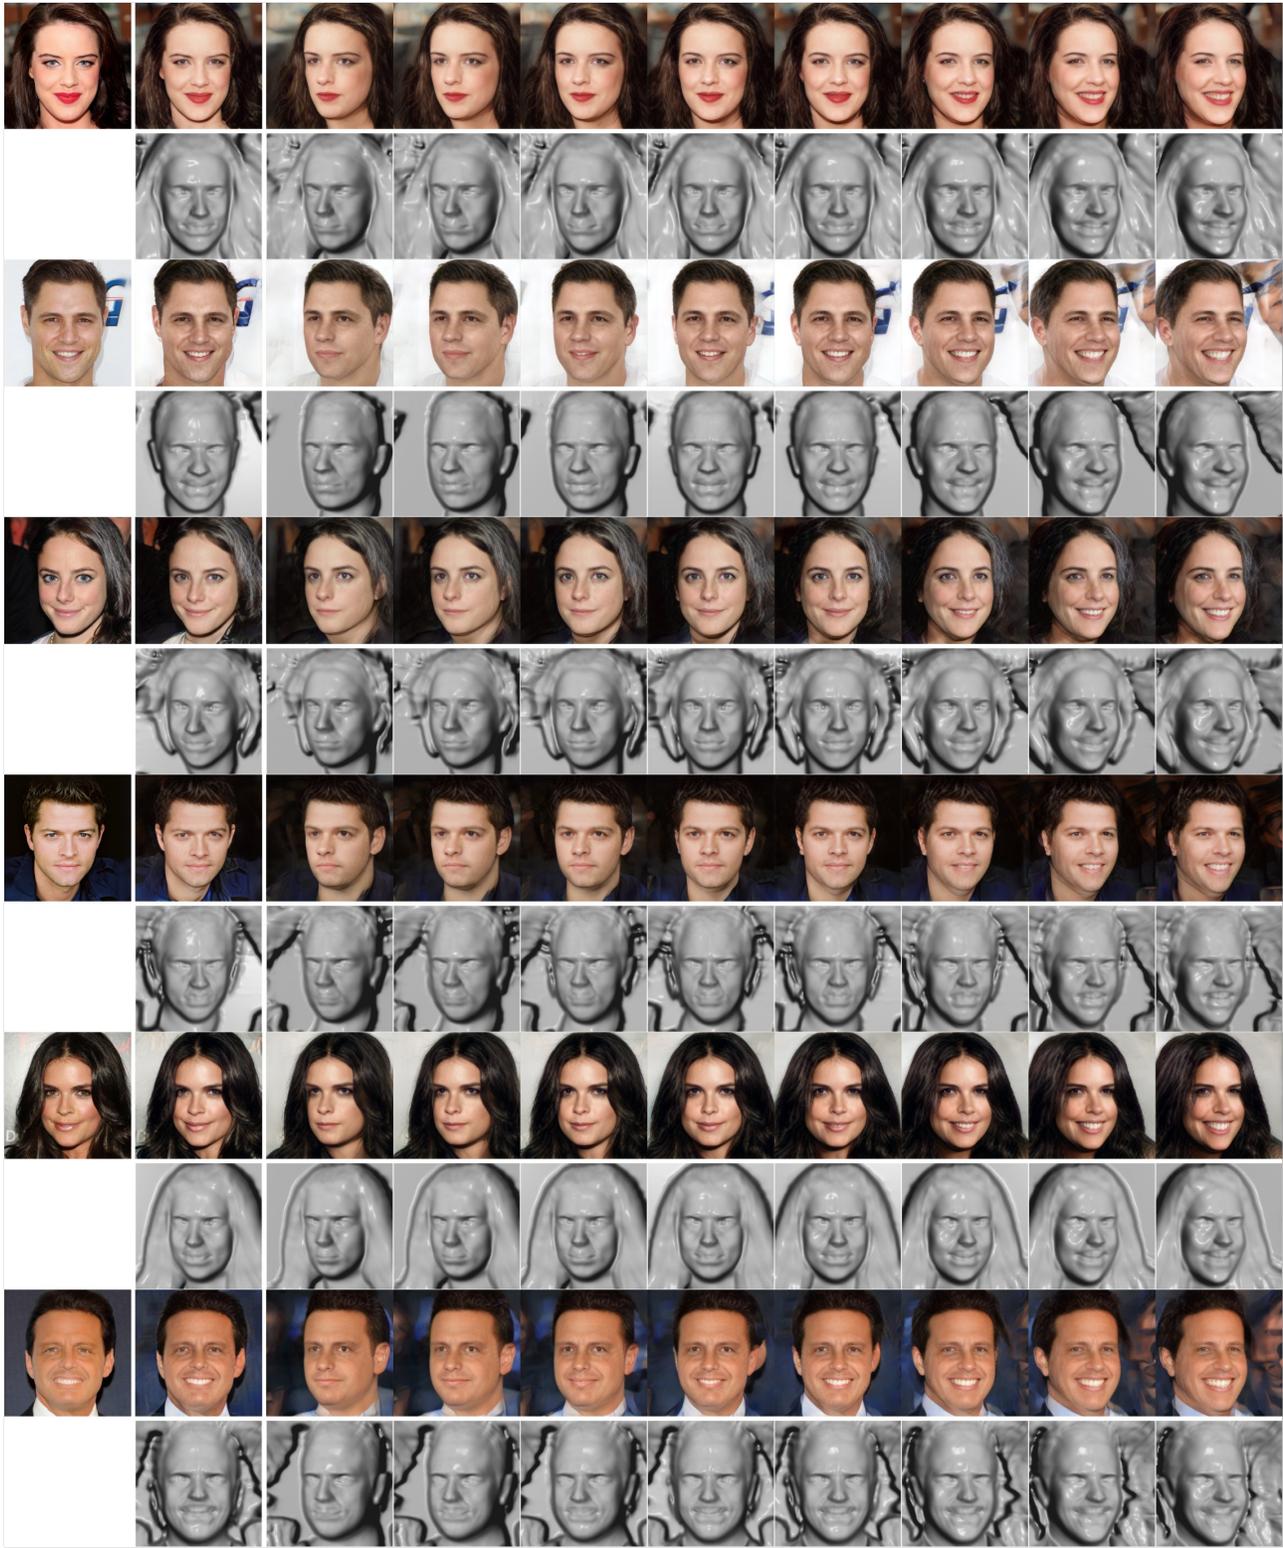


Figure 3. Visual comparisons on encoder-based methods. 'Rec' and 'Edit' represent reconstruction and editing, respectively.



Input

Ours (Rec)



Figure 4. Visual comparisons on face editing (Smile).



Input

Ours (Rec)

+ Beard / Hair

Figure 5. Visual comparisons on face editing (Beard / Hair).



Input

Ours (Rec)

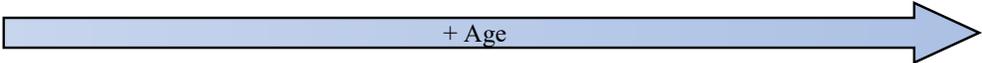


Figure 6. Visual comparisons on face editing (Age).



Input

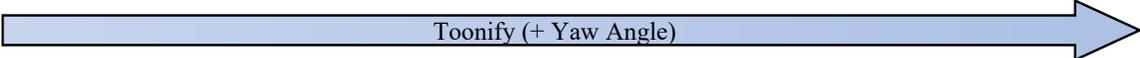
Ours (Rec)



Figure 7. Visual comparisons on face editing (Bangs).



Input



Toonify (+ Yaw Angle)

Figure 8. Toonify results on faces.

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*, 2019. 3
- [2] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and G. Wetzstein. Pi-GAN: Periodic Implicit Generative Adversarial Networks for 3D-Aware Image Synthesis. In *CVPR*, 2021. 3
- [3] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022. 1, 3
- [4] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 3
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3
- [6] Yueqi Duan, Haidong Zhu, He Wang, Li Yi, Ram Nevatia, and Leonidas J. Guibas. Curriculum DeepSDF, 2020. 2
- [7] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. StyleNeRF: A style-based 3D-aware generator for high-resolution image synthesis. In *ICLR*, 2021. 1
- [8] Yuming Jiang, Ziqi Huang, Xingang Pan, Chen Change Loy, and Ziwei Liu. Talk-to-Edit: Fine-grained facial editing via dialog. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [9] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 3, 5
- [10] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. MaskGAN: Towards diverse and interactive facial image manipulation. In *CVPR*, 2020. 5
- [11] Lars M. Mescheder, Andreas Geiger, and S. Nowozin. Which training methods for GANs do actually converge? In *ICML*, 2018. 3
- [12] 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 *ECCV*. Springer, 2020. 1
- [13] 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 *CVPR*, 2021. 1
- [14] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer. In *AAAI*, volume 32, 2018. 1
- [15] Justin NM Pinkney and Doron Adler. Resolution dependent gan interpolation for controllable image synthesis between domains. *arXiv preprint arXiv:2010.05334*, 2020. 3
- [16] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *ICLR*, 2016. 2
- [17] 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 *CVPR*, 2021. 2
- [18] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Trans. Graph.*, 2021. 3, 5
- [19] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 1
- [20] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael Black. Learning to regress 3D face shape and expression from an image without 3D supervision. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2, 5
- [21] Sefik Ilkin Serengil and Alper Ozpinar. Hyperextended lightface: A facial attribute analysis framework. In *2021 International Conference on Engineering and Emerging Technologies (ICEET)*, pages 1–4. IEEE, 2021. 2
- [22] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. InterFaceGAN: Interpreting the disentangled face representation learned by GANs. *PAMI*, PP, 2020. 2, 3
- [23] 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
- [24] Tengfei Wang, Yong Zhang, Yanbo Fan, Jue Wang, and Qifeng Chen. High-Fidelity GAN inversion for image attribute editing. In *CVPR*, 2022. 2
- [25] Xintao Wang. facexlib. <https://github.com/xinntao/facexlib>, 2020. 2
- [26] Xinyao Wang, Liefeng Bo, and Li Fuxin. Adaptive wing loss for robust face alignment via heatmap regression. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [27] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised Learning of Probably Symmetric Deformable 3D Objects from Images in the Wild. In *CVPR*, 2020. 2
- [28] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. PixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. 1, 2
- [29] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 2