

NaviNeRF: NeRF-based 3D Representation Disentanglement by Latent Semantic Navigation

Bao Xie^{1,2,3} Bohan Li^{2,4} Zequn Zhang^{2,5} Junting Dong⁶ Xin Jin^{2,3,*} Jingyu Yang¹ Wenjun Zeng^{2,3}
¹Tianjin University ²Eastern Institute of Technology, Ningbo ³Ningbo Institute of Digital Twin
⁴Shanghai Jiao Tong University ⁵Northwest Normal University ⁶Zhejiang University

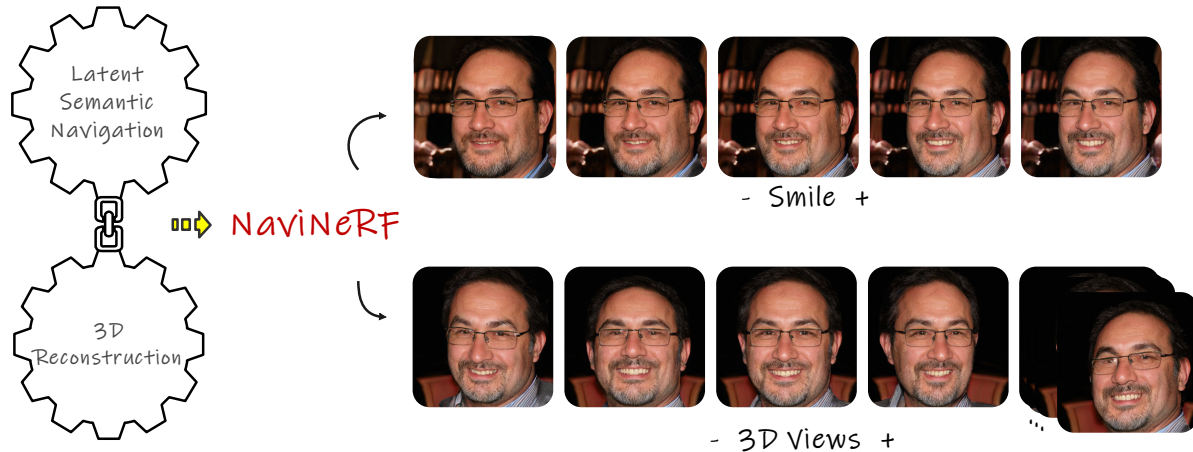


Figure 1. Generated 3D objects by NaviNeRF – a model aims to achieve fine-grained 3D disentanglement by bridging 3D reconstruction and latent semantic manipulation. The top row presents the results of shifting along the learned semantic direction that represents continuous changes in a man’s mouth, visually looks like a “smile” expression. The bottom row showcases the results of multi-view generation, which demonstrates that the attribute manipulation could still remain consistent across different views.

Abstract

3D representation disentanglement aims to identify, decompose, and manipulate the underlying explanatory factors of 3D data, which helps AI fundamentally understand our 3D world. This task is currently under-explored and poses great challenges: (i) the 3D representations are complex and in general contains much more information than 2D image; (ii) many 3D representations are not well suited for gradient-based optimization, let alone disentanglement. To address these challenges, we use **NeRF** as a differentiable 3D representation, and introduce a self-supervised **Navigation** to identify interpretable semantic directions in the latent space. To our best knowledge, this novel method, dubbed **NaviNeRF**, is the first work to achieve fine-grained 3D disentanglement without any priors or supervisions. Specifically, NaviNeRF is built upon the generative NeRF pipeline, and equipped with an **Outer Navigation Branch** and an **Inner Refinement Branch**. They are complementary — the outer navigation is to identify global-view semantic directions, and the inner refinement dedicates to fine-grained attributes. A synergistic loss is further devised

to coordinate two branches. Extensive experiments demonstrate that NaviNeRF has a superior fine-grained 3D disentanglement ability than the previous 3D-aware models. Its performance is also comparable to editing-oriented models relying on semantic or geometry priors.*

1. Introduction

3D reconstruction aims to create a virtual representation of an object or scene based on point cloud, voxel, 3D mesh, and etc. Despite significant progress of explicit reconstruction technologies such as Structure from Motion (SfM) [1], Multi-View Stereo (MVS) [2] and Structured Light (SL) [3], it remains a critical problem that the reconstructed scenes typically lack interpretability and controllability. Thus, it is important to study the 3D representation disentanglement, in which we can identify, decompose, and manipulate the underlying explanatory factors hidden in the observed 3D data.

However, 3D representation disentanglement is currently

*Denotes the corresponding author. Code is available at [this link](#).

under-explored and faces great challenges: On one hand, the 3D representations are complex with the prohibitive storage costs, which in general contains much more information than 2D image, like depth, viewpoint, etc. On the other hand, many high-dimensional 3D representations (e.g., discrete point cloud, mesh, voxel) are essentially not well suited for gradient-based optimization [4], which further increases the difficulty of disentanglement. All in all, how to efficiently and effectively achieve fine-grained 3D disentanglement without extra auxiliary priors or supervisions urgently needs to be solved.

Recently, the development of implicit representation learning has significantly promoted 3D reconstruction w.r.t the model flexibility and generalizability [5, 6, 7]. As a landmark of implicit 3D reconstruction, Neural Radiance Fields (NeRF) [4] maps scenes into a multi-layer perceptron (MLP) from limited views, which results in accurate, efficient, and differentiable 3D representations. Moreover, as a deep neural model, NeRF has preliminary shown its capability w.r.t disentangled representation learning in few studies [8, 9, 10]. Typically, the conditional NeRFs [11, 12] achieve disentanglement with pre-defined extra latent codes, which inevitably limits the diversity of decomposed attributes. On the other branch, the editing-oriented NeRFs [13, 14] also achieve a controllable 3D synthesis. However, these approaches heavily relied on geometric priors and did not identify the underlying semantic representation, such specific priors largely limited the scope of practical applications.

Revisit the recent success of disentanglement in 2D image, we knew that traversing semantically meaningful directions in the Generative Adversarial Network’s (GAN) [15] latent space leads to coherent variations in the generated 2D image [16, 17, 18, 19]. Typically, the way of smooth navigation [20, 21] is investigated for GAN-based semantic editing in the space of the generator’s parameters. These observations indicate that the underlying explanatory properties are probably embedded in the generative latent space.

Based on the above discussions, in this paper, **we explore to use NeRF as a differentiable 3D representation, and introduce a self-supervised navigation to identify interpretable semantic directions in the generative latent space.** We name this novel method as **NaviNeRF**. As shown in Figure 1, NaviNeRF achieves a fine-grained 3D disentanglement by bridging 3D reconstruction and latent semantic manipulation. When shifting along the disentangled semantic direction that represents the mouth, we obtain a group of continuously changed visual results, look like a “smile”. In addition, the generated results of NaviNeRF could remain 3D consistency well across different views.

Specifically, Figure 2 showcases that NaviNeRF is composed of two main components: an outer navigation branch and an inner refinement branch. The outer navigation aims

to identify the traversal directions as global-view factors in the latent space for disentangled representation learning — this process employs a learnable matrix to append a shift on a latent code. The shifted code, paired with the original one, are used to generate a pair of images through the pre-trained generator. A trainable decoder is then devised to **predict the shift (i.e., semantic direction)** based on such paired images, with a reconstruction loss [20]. Similarly, the inner branch dedicates to more fine-grained attributes by appending shifts on the *specific* dimensions of intermediate latent code. Finally, a synergistic loss function is further designed to combine these two complementary branches well. Compared to off-the-shelf solutions, NaviNeRF does not resort to explicit conditional codes or any geometry priors. In summary, our contributions are:

1. To our best knowledge, the proposed NaviNeRF is the first work that could achieve fine-grained 3D disentanglement at feature-level, without any priors and additional supervision.
2. We take full advantage of both latent semantic navigation (the outer branch) and NeRF representation (the inner branch) in a complementary way. The outer branch learns to identify semantic directions for global disentangled representation learning, and the inner branch learns to focus on fine-grained attributes.
3. As a by-product, a simple synergistic loss is designed to collaborate well two outer-inner branches within NaviNeRF.

We evaluate NaviNeRF on two popular benchmarks: FFHQ [22] for the human face and AFHQ [23] for the animal face. NaviNeRF outperforms typical 3D-aware GANs including pi-GAN [24], GIRAFFE [25] and StyleNeRF [26] in attribute manipulation. Furthermore, the model obtains comparable performance to editing-oriented models which rely on semantic or geometric priors. Extensive ablation studies are also conducted to support our claims.

2. Related Works

Our work strongly relies on NeRF, GAN and theory of the latent semantic analysis. This section describes the relevant studies in these areas.

Neural implicit representations. Encoding a scene into neural networks has shown great promise as a research direction. This includes, but is not limited to: parameterizing the geometry of a scene via signed distance functions or occupancy field [27, 28], encoding both geometry and appearance [4], etc. One notable contribution in this field, is the work known as NeRF that has drawn extensive attention recently. It encodes a scene as 5D vector-valued function approximated by a MLP, where the function denotes

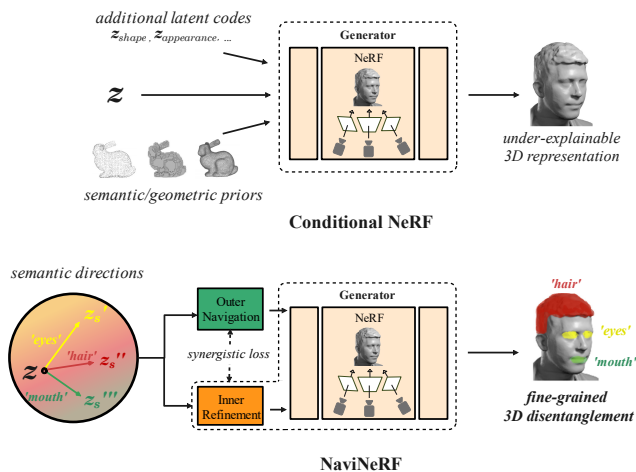


Figure 2. Workflows of standard conditional NeRFs and NaviNeRF. NaviNeRF combines an outer navigation branch and an inner refinement branch by a synergistic loss, for fine-grained 3D disentanglement. Compared with existing solutions, NaviNeRF does not require conditional latent codes or semantic/geometric priors.

spatial coordinates, viewing direction, color and density. This network can be repeatedly called by any of the volume rendering techniques to produce novel views [29, 30]. The impressive performance of NeRF inspired follow-up works to extend it in alternative settings, such as training from unstructured images [31, 32], training without camera poses [33, 29, 34], training with generative models [35, 36], etc. As a differentiable representation, NeRF and its variants have demonstrated strong capabilities at generating 3D scenes with high accuracy, efficiency and consistency. However, these approaches are commonly deficient in interpretable control over partial properties, such as shape, color, texture, and lighting.

3D-aware GAN models. Recently, generative models have brought NeRF a certain degree of scene control capabilities. The early attempts in this routine are GRAF [37] and pi-GAN [24]. The former handles category-specific generation by conditioning NeRF on shape and appearance codes. Following the NeRF pipeline, the generator can synthesize an image by taking random codes and camera poses. The generated image is fed into the discriminator along with real images, thus implementing a GAN. pi-GAN is similar to GRAF, but conditions on a single latent code and utilizes FiLM-SIREN layers [38, 39] instead of simple MLPs. Encoding additional latent codes enriches the model with disentangled capability, but double-edgedly, limits the range of disentangled attributes.

Another impressive work in this domain is GIRAFFE [25], which represents scenes as compositional generative NeRFs without any additional supervision. It is

in capacity of disentanglement on separating background and foreground of the scene. Although the compositional architecture delivers control capacity over object-level, individual’s local attributes are not yet fully disentangled. More recently, [26] integrates NeRF into a StyleGAN [40] based generator to produce high-resolution and multi-view consistent 3D scenes. As a by-product, it inherits the style control ability from StyleGAN baseline but still, fails to disentangle on detailed attributes.

3D representation editing methods. Except embedding additional codes into latent space, many editing-oriented NeRFs adopt option to enhance 3D perception by leveraging semantic or geometric priors [13, 41, 42]. With GAN inversion technology [17, 43, 44], these approaches can edit specific regions of a 3D scene under interactive controls given by user. Typically, [45] trains the model using paired monocular images and semantic maps, and obtains locally-editable images. Furthermore, [46] extracts an explicit triangular mesh representation as geometric priors, which can then be intuitively deformed by the user for 3D editing. Although these methods could obtain promising results on pixel-wise editing, they are still inapplicable to perform interpretable disentanglement since the underlying semantic representations are not essentially learned.

Latent semantic analysis. Empirical studies have revealed that GAN latent spaces are embedded with interpretable semantic information [16, 47, 48, 49]. Representatively, [50] verifies that GANs trained with face images have latent spaces that contain semantic directions corresponding to specific facial features. Since such interpretable directions provide a straightforward route to robust image editing, their discovery currently receives much research attention. [48] further proposes the classifiers pre-trained on the facial data, to predict certain face attributes. These classifiers are then used to produce pseudo-labels for the generated images and their latent codes. Furthermore, [17] conducts series of experiments to verify that the interpretable directions are responsible for diverse specific features, by maximizing the score of pre-trained generative models. The aforementioned semantic analysis could potentially enrich the exploration of latent space in generative NeRFs, which therefore resolve the defectiveness of semantic information in 3D scenarios. Motivated by this idea, our model is proposed to complementally take advantage of NeRF representation and semantic direction manipulation, targeting on fine-grained 3D disentanglement.

3. Methodology

In this section, we elaborate the three key modules that constitute our model: an outer navigation branch, an inner refinement branch and a synergistic loss for interaction. The architecture of the model is illustrated in Figure 4. Correspondingly, we start the introduction of the navi-




Figure 3. Within a latent space \mathcal{Z} , the model is proposed to discover and manipulate an interpretable semantic direction from original code z to shifted code z_s . Traversing along this direction lead to continue changes on a disentangled representation of generated image.

gation branch in Section 3.1, and elaborate the structure of the refinement branch in Section 3.2. Then in Section 3.3, the synergistic loss coordinating the two branches are described, along with other loss functions we employed.

3.1. Outer Navigation Branch

The navigation branch, inspired by previous studies of latent semantic manipulation [20, 21, 51], is proposed to identify interpretable semantic directions in the latent space. Specifically, given a pre-trained GAN generator G , which maps the samples $z \sim N(0, \mathbb{I})$ drawn from the latent space \mathcal{Z} to generate multi-view images. Our objective is to learn a set of semantic directions, that facilitate continuous changes of specific attributes in the generated image upon being shifted along each direction. For example, Figure 3 visualizes the alter on the object’s whiskers in RGB space, which aroused by the traverse from z to z_s in latent space.

Towards this target, we propose a learnable matrix $S \in \mathbb{R}^d$ that enables the identification of a shift with a specific direction index β and scale ϵ . Here, d represents the dimensionality of the matrix, where the columns of S corresponding to candidate directions. The shift S is appended on z for conducting a shifted code z_s . Both z and z_s are then input to the pre-trained G to produce paired images I_0 and I_1 respectively:

$$I_0 = G(z), I_1 = G(S(\epsilon \cdot z\beta))$$

where the diversity between the two images being solely attributed to this latent shift. In other words, I_1 is the transformation of I_0 , corresponding to moving by ϵ along the direction β .

After the reconstruction, a trainable decoder D is proposed to project generated images from RGB space back to latent space. More specifically, D is a function that maps images pairs into the shift increment, parameterized by a MLP network. The shift increment $(\Delta\beta', \Delta\epsilon')$ from I_0 to

I_1 , is constrained with the ground truth by the reconstruction loss $\mathcal{L}_{\mathcal{R}}$.

Although the outer module is devised to discover semantic directions, it alone is insufficient for 3D scenarios owing to the deficiency of geometric consistency. Nevertheless, additional instructions are expected, for model to concentrate on fine-grained attributes over discovered directions. Towards it, we propose an inner refinement branch for two goals: obtaining the perception of fine-grained representations and, preserving 3D consistency.

3.2. Inner Refinement Branch

Within the inner branch, we adopt StyleNeRF baseline as the generator, which takes integrated advantages of NeRF and StyleGAN. In a basic GAN, latent code z is sampled directly from a Gaussian distribution and determines the global style of the generated image. However, the limited capacity of the normal distribution constrains the disentanglement capability of \mathcal{Z} [52]. Differently, StyleGAN maps native z to a layer-wise style code w by a 8-layer mapping network M . The intermediate latent space is referred as \mathcal{W}^+ space that contains more disentangled features than \mathcal{Z} . Magnetized by the disentanglement capability of \mathcal{W}^+ , we devise to build paired codes (original and shifted) similar as the outer branch, but with the shifts over w .

Tentatively, semantic shifts are appended on each dimension of latent code w , which however, leads to an unexpected entanglement among global style and fine-grained details. To our knowledge, the phenomenon is attributed to the different control scope of distinct dimensions in w . Specifically, it has been observed that dimensions of w correspond to different levels of details, roughly in three groups: global, coarse, and fine [40, 53, 54, 55]. Conditioning partial details together with global and coarse style can pose a challenge for disentanglement. In this sense, we turn to align the shifts on 9th - 18th dimensions of w which theoretically, controls the fine-grained attributes, therefore encouraging the model to learn partial representations.

Following the architecture of StyleNeRF, we perform NeRF++ behaving as the NeRF synthesis network. It comprises a foreground (FG) NeRF in a unit sphere and a background (BG) NeRF represented using an inverted sphere parameterization [56]. As shown in Figure 4, two MLPs are utilized separately to predict the density. A shared MLP is then employed with up-sample blocks for color prediction.

To achieve disentangled semantic manipulation, the shifted dimensions (9th - 18th) of w are two-to-one fed into 5th - 9th NeRF MLP layers through an affine transformation, followed by a decoder to predict shift increment $(\Delta\beta'', \Delta\epsilon'')$.




Figure 4. NaviNeRF is characterized by two complementary branches, termed as the outer navigation branch and the inner refinement branch. The former, depicted in green, appends a shift on sampled latent code z through a learnable matrix S . z and shifted code z_s are used to generate paired images, which are devoted to train the decoder D for semantic direction identification. The latter, shown in orange, produces fine-grained awareness and 3D consistency by appending shifts on specific dimensions of intermediate latent vector w . The two branches are combined by a synergistic loss, ultimately achieving feature-level 3D disentanglement.

3.3. Loss Functions

The heart of the NaviNeRF is at the complementarity of navigation and refinement branches. To combine these two modules, a synergistic loss is devised in three steps: (i) decoding generated image pairs into the shift with the direction index β and scale ϵ ; (ii) calculating the increments of both branches, termed as $(\Delta\beta', \Delta\epsilon')$ and $(\Delta\beta'', \Delta\epsilon'')$ respectively; (iii) minimizing the distance of two increments. The variation of the two modules can be calculated as:

$$(\Delta\beta', \Delta\epsilon') = D(G(z), G(S(\epsilon \cdot z\beta)))$$

$$(\Delta\beta'', \Delta\epsilon'') = D(G(z), G(\epsilon \cdot w\beta))$$

Generally, the synergistic loss is proposed to be a cross-entropy loss, which therefore can be demonstrated as:

$$\mathcal{L}_S = \text{CrossEntropy}((\Delta\beta', \Delta\epsilon'), (\Delta\beta'', \Delta\epsilon''))$$

In addition to the synergistic loss, we utilize the reconstruction loss \mathcal{L}_R on generated images and ground truth to optimize the reconstruction quality of outer branch using a MSE loss.

$$\mathcal{L}_R = \text{MSE}((\Delta\beta', \Delta\epsilon'), (\beta, \epsilon))$$

We also apply the consistency loss \mathcal{L}_C to enforce 3D consistency, as instituted by [26]. More formally, we propose another original NeRF path without up-sampling blocks in the NeRF generator, for producing a low-resolution but consistent image to supervise I_2 by the consistency loss. In this

way, I_2 can be closer to the NeRF results, which have multi-view consistency. The loss is calculated as follows:

$$\mathcal{L}_C = \frac{1}{|P|} \sum_{(i,j) \in P} (I_2[i, j] - I_{\text{NeRF}}[i, j])^2$$

where P denotes randomly sampled pixels. I_2 is instantiated as the image generated from the inner branch and I_{NeRF} is the image from the original NeRF.

In such setting, the total loss can be summarized as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_S + \lambda_R \mathcal{L}_R + \lambda_C \mathcal{L}_C$$

where the λ_R and λ_C are the hyper-parameter. In default, we adopt $\lambda_R = 0.8$ and $\lambda_C = 0.6$ to balance the disentanglement capability and reconstruction quality.

4. Experiments

4.1. Experimental Settings

We conduct a set of experiments on several datasets: **FFHQ** [22] consists of 70,000 high-quality images of human faces; **AFHQ** [23] contains 15,000 high-quality images at a resolution of in three categories of cat, dog, and wildlife; **CompCars** [57] contains 136726 images capturing the entire cars with different styles; **LSUN** [58] consists of about one million images for multiple object categories. We pre-train the generator with resized images from aforementioned datasets at 256×256 resolution for a trade-off on quality and controllability. For pre-training, we follow the instructions outlined in StyleNeRF adopting batch size in 64 and a learning rate as 0.0025. To train outer and inner branches, we apply batch sizes of 64 and 32 for FFHQ and




Figure 5. Fine-grained 3D Disentanglement Results of NaviNeRF. The left columns present the results of attribute manipulation and the right columns showcase corresponding 3D reconstruction results. Respectively, (a) demonstrates the semantic manipulation on the FFHQ dataset including the man’s mouth, whiskers and the girl’s hair; (b) shows the manipulation results on the puppy’s ears, tongue and cheeks.

AFHQ, respectively, with a learning rate of 0.0005. All experiments were performed on 4 Nvidia GPUs (Tesla A100 80GB) with CUDA version 11.6.

4.2. Results

4.2.1 Fine-grained 3D Disentanglement

We firstly conduct experiments on two popular datasets: FFHQ and AFHQ, in aim to evaluate the effectiveness of our proposed model in feature-level 3D disentanglement. These benchmarks consist of high-resolution images of single objects and cover a diverse range of real-world scenarios, highlighting the robustness of our results. We further expand the model’s applicability to more intricate scenes by leveraging CompCars and LSUN datasets, which effectively showcasing the model’s capacity for generalization.

In Figure 5, we present several visual effects on different objects induced by the discovered directions, paired with their reconstruction results. More specifically, we illustrate the attribute manipulation results on the left side. And the results of the 3D reconstruction are demonstrated on the right, visually formed as the multiple perspectives of the image after manipulation (the 5th image on the left). It is worth mentioning that the quality of the reconstructed images is

highly dependent on the pre-trained generator. But even the model is fully configured as StyleNeRF in the pre-training stage, artifacts still exist in generated images for the AFHQ dataset. However, with such ambiguous and challenging data, the model still achieves fine-grained awareness, which validates its robustness.

Different from those pixel-level editing works, our model achieves semantic manipulation in latent space, which is easily extendable to more scenes per the theory of disentangled representation learning [59]. Moreover, the differentiable and continuous nature of the NeRF-based implicit representation makes it more flexible and geometry-free for novel view synthesis. Furthermore, NaviNeRF also inherits the well-studied properties of StyleGAN, leading to a universal feature generation capability. To emphasize the generalization capability of the model, we showcase extra fine-grained disentanglement results within more generic scenes in Figure 6. Supplementary examples can be found in the appendix and in the repository.

4.2.2 Qualitative Comparison

NaviNeRF vs. typical 3D-aware GANs. Typical 3D-aware GAN such as pi-GAN, GIRAFFE and StyleNeRF




Figure 6. Fine-grained 3D disentanglement results on LSUN and CompCars datasets that covering more general scenes, to demonstrate both the disentanglement and generalization capability of the model.

claim a certain degree of disentanglement in their model. Thereinto, pi-GAN and GIRAFFE provide scene control over object appearance, style, and rotation by altering shape code z_{shape} and appearance code z_{app} . On the other hand, StyleNeRF allows for the manipulation of global styles by leveraging the disentanglement capacity of StyleGAN. Since GIRAFFE and StyleNeRF provide pre-trained models on the FFHQ, we directly load their checkpoints and re-trained pi-GAN with the same data. However, as these models control the scenes in different ways, comparing and visualizing the results of 3D semantic manipulation can be a challenging task. To ensure fairness, we adopt style interpolation for each typical GAN over the same attribute to demonstrate their capability of fine-grained control. Specifically as shown in Figure 7, we manually select two images to determine an assumed direction for these 3D-aware GANs without the function of direction manipulation. In the first image, the mouth of the object is closed while in the second image, the mouth is smiling. The other regions of the face in both images are identical. We then extract the paired latent code z and z_s corresponding to these two images. By further manipulating along the direction from z to z_s , we can thus conduct specific attribute (mouth) control results for each model.

In the first three rows of Figure 7, we represent the comparison results of continue interpolation upon the man’s mouth. The results obtained by pi-GAN and GIRAFFE show that global styles such as beard, skin tone and eye socket are changing simultaneously during the manipulation on the mouth. Although StyleNeRF produces better

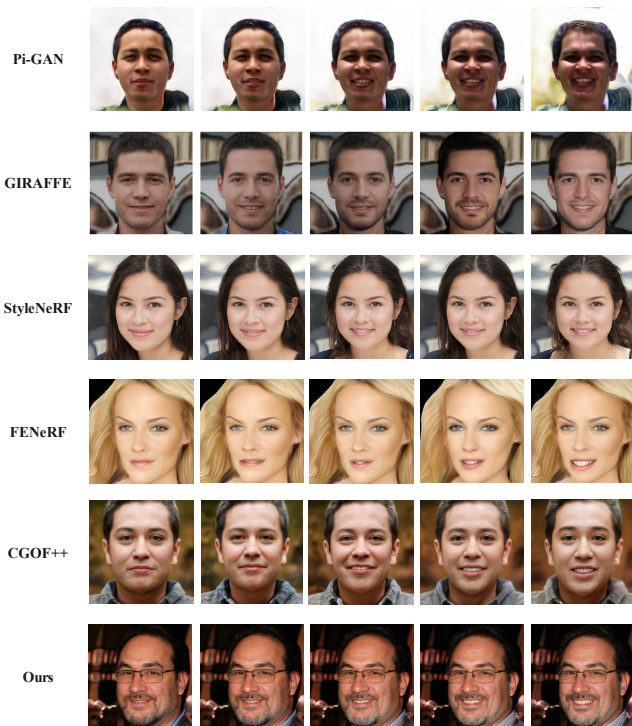


Figure 7. Comparison results of NaviNeRF with typical 3D-aware GANs and 3D editing-oriented models on the FFHQ dataset. The first three rows present disentanglement results on manipulating object’s mouth from pi-GAN, GIRAFFE and StyleNeRF, respectively. The 4th - 5th rows demonstrate the manipulation results of two editing-oriented models.

results, some partial representations remain entangled such as hairstyle. NaviNeRF demonstrates better overall performance on disentangling specific attributes compared to other baselines, which confirms the superiority of the model for feature-level 3D disentanglement.

NaviNeRF vs. 3D editing-oriented models. Alternately, another bunch of work tends to attain 3D scene control by incorporating semantic or geometric priors. Encoded priors can enrich the model 3D perception and enable editing of specific attributes through inversion techniques. Accordingly, we compare NaviNeRF with two novel works reported in 2022 within this domain: FENeRF [45] and CGOF++ [13]. The former learns 3D representation from widely available monocular images and semantic mask pairs. The model uses the semantic mask to manipulate partial attributes via GAN inversion. The latter is a conditional NeRF that incorporates a mesh-guided sampling process and a depth-aware density regularizer. For comparison, we load their pre-trained models on the FFHQ dataset and conduct the samples for editing object’s mouth.

The last three rows of Figure 7 demonstrate compari-

son results for FENeRF, CGOF++, and NaviNeRF. Three methods achieve approximate results on editing the partial attribute. Although the extra priors offer a shortcut for scene disentanglement, the editing-oriented models inevitably suffer from its redundancy, inflexibility and inefficiency. In addition, these models are dictated to edit a specific pixel area but not essentially understand the underlying semantic meanings. As emphasized earlier, our model aims to autonomously learn the latent semantic information, without any priors and additional supervision. Therefore, the comparable results demonstrate the superiority of our model.

4.2.3 Quantitative Comparison

Table 1 reports the results of Frechet Inception Distance (FID) [60] and Kernal Inception Distance (KID) [60] scores to measure the quality of the generated images. Our model outperforms other typical 3D-aware models and is slightly inferior to StyleNeRF. To our understanding, the decrease in generation quality compared with StyleNeRF can be a trade-off for fine-grained control (i.e., we did not fine-tune the pre-trained generator in attribute manipulation). Furthermore, the performance decrease in the AFHQ dataset, which we believe, is caused by the non-fully trained generator as noted in Section 4.2.1.

Table 1. Quantitative comparison results with typical 3D-aware models in FID and KID $\times 10^3$.

Models	FFHQ		AFHQ	
	FID	KID	FID	KID
pi-GAN	87	99	53	35.4
GIRAFFE	38	25.7	36	14.7
StyleNeRF	10.4	4.6	16	4.3
Ours	13	6.9	22	9.1

Figure 8 demonstrates the FID scores for images with various shifting magnitudes in the “baldness” direction. The results indicate that the model can maintain the generation quality across the attribute manipulation. It supports our claims that the model has learned to manipulate fine-grained features while maintaining a consistent global style.

Table 2 and 3 presents the method budget on the FFHQ to reveal that NaviNeRF is high-efficient. Table 2 shows that two-branch design did not increase the feature size and we also use the pre-trained generator for lower cost, thus avoid high memory issue in rendering. As demonstrated in Table 3, NaviNeRF achieves significant speedup over pure NeRF-based methods and is comparable to StyleNeRF.




Figure 8. FID scores of images with different shifting magnitudes on the man’s baldness.

Table 2. Training Budget report on parameters size, FLOPs, memory cost and training time (TT).

Budget	Params(M)	FLOPs(G)	Mem(G)	TT(hrs)
Single Branch	11.9	142.9	23	5.8
Double Branches	12.2	162.4	25	7.5

Table 3. Rendering time (RT) comparisons with other methods at 256^2 .

Models	pi-GAN	GIRAFFE	StyleNeRF	Ours
RT(ms/image)	785	181	75	97

4.3. Ablation study

To validate the effectiveness of some key designs in NaviNeRF, we conduct ablations over w/o synergistic loss, choice of shifting dimensions and choice of latent space.

Shifting on w dimensions. As previously mentioned, different w dimensions correspond to three levels of partial style: 1st - 4th layers determine the global features, 5th - 8th layers for coarse features and 9th - 18th layers for fine-grained details. In Figure 9, we compare the reconstruction results of shifting on 9th - 18th layers (full model) against with shifting on every 18 layer (second row) in w . When shifting on every w dimension, global features such as face shape, skin texture and wrinkle style are variously entangled, resulting in an older appearance for the girl. It indicates that shifting on specific dimensions can enforce the model to concentrate on fine-grained representations.

\mathcal{W} space vs. \mathcal{W}^+ space. Referring to StyleNeRF, \mathcal{W} is an intermediate space with a distribution matching better to the real data compared with the original \mathcal{Z} space. It contains a single intermediate w vector, whereas the \mathcal{W}^+ space




Figure 9. Ablation studies on manipulating the object’s “hair” over the choice of shifting dimensions, the choice of latent space and w/o synergistic loss.

comprises 18 different style vectors. To examine the impact of latent space choice, we append identical shifts on \mathcal{W} and \mathcal{W}^+ . In the case of \mathcal{W} space, we duplicate the single w into 18 dimensions. The semantic shifts are then appended on the 9th - 18th dimensions of the intermediate vectors from both spaces. The third row in Figure 9 demonstrates that, shifting on \mathcal{W} space enable the model to control global styles but fails to achieve fine-grained manipulation. That is why we choose to propose our model on \mathcal{W}^+ for feature-level 3D disentanglement.

w/o synergistic loss. We perform NaviNeRF with or without the synergistic loss to study its impact. Since the inner branch cannot be directly trained without the synergistic loss, we replace it with a reconstruction loss which is in the same configuration as $\mathcal{L}_{\mathcal{R}}$. The 4th row of Figure 9 demonstrates that removing the synergistic loss leads to severe 3D inconsistent artifacts during disentanglement. The result supports the notion that synergistic loss combines the outer and inner branches in a complementary way to achieve fine-grained disentanglement and 3D consistency.

5. Conclusion

In this paper, we present NaviNeRF, a NeRF-based 3D reconstruction model that achieves fine-grained disentanglement while preserving 3D accuracy and consistency without any priors and supervision. The model consists of two complementary branches: an outer navigation branch delicate to identify the traversal directions as factors in the latent space, while an inner refinement branch produces fine-grained awareness and 3D consistency. We also design

a synergistic loss to combine the two modules. The model is evaluated on challenging datasets to demonstrate its ability of fine-grained disentanglement in 3D scenarios. The experimental results indicate that NaviNeRF outperforms typical conditional NeRFs. Furthermore, its performance is also comparable to editing-oriented models relying on semantic or geometry priors, which supports our claims.

Acknowledgement

This work was supported in part by ZJNSFC under Grant LQ23F010008.

References

- [1] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 1
- [2] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015. 1
- [3] Andrew Forbes, Michael de Oliveira, and Mark R Dennis. Structured light. *Nature Photonics*, 15(4):253–262, 2021. 1
- [4] 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. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [5] Alireza Khatamian and Hamid R Arabnia. Survey on 3d surface reconstruction. *Journal of Information Processing Systems*, 12(3):338–357, 2016. 2
- [6] George Fahim, Khalid Amin, and Sameh Zarif. Single-view 3d reconstruction: A survey of deep learning methods. *Computers & Graphics*, 94:164–190, 2021. 2
- [7] Siming Yan, Zhenpei Yang, Haoxiang Li, Li Guan, Hao Kang, Gang Hua, and Qixing Huang. Implicit autoencoder for point cloud self-supervised representation learning. *arXiv preprint arXiv:2201.00785*, 2022. 2
- [8] Wobong Jang and Lourdes Agapito. Codenerf: Disentangled neural radiance fields for object categories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12949–12958, 2021. 2
- [9] Ziyu Wang, Yu Deng, Jiaolong Yang, Jingyi Yu, and Xin Tong. Generative deformable radiance fields for disentangled image synthesis of topology-varying objects. *arXiv preprint arXiv:2209.04183*, 2022. 2
- [10] Shunyu Yao, RuiZhe Zhong, Yichao Yan, Guangtao Zhai, and Xiaokang Yang. Dfa-nerf: personalized talking head generation via disentangled face attributes neural rendering. *arXiv preprint arXiv:2201.00791*, 2022. 2
- [11] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3835–3844, 2022. 2

- [12] Kyungmin Jo, Gyumin Shim, Sanghun Jung, Soyoung Yang, and Jaegul Choo. Cg-nerf: Conditional generative neural radiance fields. *arXiv preprint arXiv:2112.03517*, 2021. 2
- [13] Keqiang Sun, Shangzhe Wu, Ning Zhang, Zhaoyang Huang, Quan Wang, and Hongsheng Li. Cgof++: Controllable 3d face synthesis with conditional generative occupancy fields. *arXiv preprint arXiv:2211.13251*, 2022. 2, 3, 7
- [14] 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. *ACM Transactions on Graphics (TOG)*, 41(6):1–10, 2022. 2
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2
- [16] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. Ganalyze: Toward visual definitions of cognitive image properties. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5744–5753, 2019. 2, 3
- [17] Ali Jahanian, Lucy Chai, and Phillip Isola. On the “steerability” of generative adversarial networks. *arXiv preprint arXiv:1907.07171*, 2019. 2, 3
- [18] William Peebles, John Peebles, Jun-Yan Zhu, Alexei Efros, and Antonio Torralba. The hessian penalty: A weak prior for unsupervised disentanglement. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, pages 581–597. Springer, 2020. 2
- [19] Binxu Wang and Carlos R Ponce. The geometry of deep generative image models and its applications. *arXiv preprint arXiv:2101.06006*, 2021. 2
- [20] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. In *International conference on machine learning*, pages 9786–9796. PMLR, 2020. 2, 4
- [21] Anton Cherepkov, Andrey Voynov, and Artem Babenko. Navigating the gan parameter space for semantic image editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3671–3680, 2021. 2, 4
- [22] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 2, 5
- [23] 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. 2, 5
- [24] Eric R Chan, Marco Monteiro, Petr Kellnhöfer, 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. 2, 3
- [25] 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, 3
- [26] 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. 2, 3, 5
- [27] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. 2
- [28] Xudong Xu, Xingang Pan, Dahua Lin, and Bo Dai. Generative occupancy fields for 3d surface-aware image synthesis. *Advances in Neural Information Processing Systems*, 34:20683–20695, 2021. 2
- [29] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. 3
- [30] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 3
- [31] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 3
- [32] Takuhiko Kaneko. Ar-nerf: Unsupervised learning of depth and defocus effects from natural images with aperture rendering neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18387–18397, 2022. 3
- [33] Quan Meng, Anpei Chen, Haimin Luo, Minye Wu, Hao Su, Lan Xu, Xuming He, and Jingyi Yu. Gnerf: Gan-based neural radiance field without posed camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6351–6361, 2021. 3
- [34] Shu Chen, Yang Zhang, Yaxin Xu, and Beiji Zou. Structure-aware nerf without posed camera via epipolar constraint. *arXiv preprint arXiv:2210.00183*, 2022. 3
- [35] Shengqu Cai, Anton Obukhov, Dengxin Dai, and Luc Van Gool. Pix2nerf: Unsupervised conditional p-gan for single image to neural radiance fields translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3981–3990, 2022. 3
- [36] Adam R Kosiorek, Heiko Strathmann, Daniel Zoran, Pol Moreno, Rosalia Schneider, Sona Mokrá, and Danilo Jimenez Rezende. Nerf-vae: A geometry aware 3d scene generative model. In *International Conference on Machine Learning*, pages 5742–5752. PMLR, 2021. 3

- [37] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems*, 33:20154–20166, 2020. 3
- [38] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020. 3
- [39] Vincent Dumoulin, Ethan Perez, Nathan Schucher, Florian Strub, Harm de Vries, Aaron Courville, and Yoshua Bengio. Feature-wise transformations. *Distill*, 3(7):e11, 2018. 3
- [40] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 3, 4
- [41] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20374–20384, 2022. 3
- [42] Kacper Kania, Kwang Moo Yi, Marek Kowalski, Tomasz Trzcinski, and Andrea Tagliasacchi. Conerf: Controllable neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18623–18632, 2022. 3
- [43] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 597–613. Springer, 2016. 3
- [44] Minyoung Huh, Richard Zhang, Jun-Yan Zhu, Sylvain Paris, and Aaron Hertzmann. Transforming and projecting images into class-conditional generative networks. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 17–34. Springer, 2020. 3
- [45] Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. Fenerf: Face editing in neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7672–7682, 2022. 3, 7
- [46] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18353–18364, 2022. 3
- [47] Antoine Plumerault, Hervé Le Borgne, and Céline Hudelot. Controlling generative models with continuous factors of variations. *arXiv preprint arXiv:2001.10238*, 2020. 3
- [48] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1532–1540, 2021. 3
- [49] Valentin Khruikov, Leyla Mirvakhabova, Ivan Oseledets, and Artem Babenko. Disentangled representations from non-disentangled models. *arXiv preprint arXiv:2102.06204*, 2021. 3
- [50] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9243–9252, 2020. 3
- [51] Xuanchi Ren, Tao Yang, Yuwang Wang, and Wenjun Zeng. Learning disentangled representation by exploiting pretrained generative models: A contrastive learning view. In *International Conference on Learning Representations*, 2021. 4
- [52] Weihao Xia, Yulun Zhang, Yujun Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. Gan inversion: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 4
- [53] 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. 4
- [54] 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. 4
- [55] X Zheng, Yang Liu, P Wang, and Xin Tong. Sdf-stylegan: Implicit sdf-based stylegan for 3d shape generation. In *Computer Graphics Forum*, volume 41, pages 52–63. Wiley Online Library, 2022. 4
- [56] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 4
- [57] Linjie Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. A large-scale car dataset for fine-grained categorization and verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3973–3981, 2015. 5
- [58] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 5
- [59] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013. 6
- [60] 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. 8