

Semantic 3D-Aware Portrait Synthesis and Manipulation Based on Compositional Neural Radiance Field

Tianxiang Ma^{1,2*}, Bingchuan Li^{3*}, Qian He³, Jing Dong^{2†}, Tieniu Tan^{2,4},

¹School of Artificial Intelligence, University of Chinese Academy of Sciences

²CRIPAC & NLPR, Institute of Automation, Chinese Academy of Sciences

³ByteDance Ltd, Beijing, China

⁴Nanjing University

tianxiang.ma@cripac.ia.ac.cn, {libingchuan, heqian}@bytedance.com, {jdong, tnt}@nlpr.ia.ac.cn,

Abstract

Recently 3D-aware GAN methods with neural radiance field have developed rapidly. However, current methods model the whole image as an overall neural radiance field, which limits the partial semantic editability of synthetic results. Since NeRF renders an image pixel by pixel, it is possible to split NeRF in the spatial dimension. We propose a Compositional Neural Radiance Field (CNeRF) for semantic 3D-aware portrait synthesis and manipulation. CNeRF divides the image by semantic regions and learns an independent neural radiance field for each region, and finally fuses them and renders the complete image. Thus we can manipulate the synthesized semantic regions independently, while fixing the other parts unchanged. Furthermore, CNeRF is also designed to decouple shape and texture within each semantic region. Compared to state-of-the-art 3D-aware GAN methods, our approach enables fine-grained semantic region manipulation, while maintaining high-quality 3D-consistent synthesis. The ablation studies show the effectiveness of the structure and loss function used by our method. In addition real image inversion and cartoon portrait 3D editing experiments demonstrate the application potential of our method.

Introduction

Photo-realistic image synthesis methods based on 2D Generative Adversarial Networks, such as StyleGAN (Karras, Laine, and Aila 2019; Karras et al. 2020), have achieved widespread success. However, they are designed for single-view image generation and do not have the ability to generate 3D-consistent images.

The 3D-aware GAN methods aim to synthesize 3D-consistent images. Many of these methods (Sitzmann, Zollhöfer, and Wetzstein 2019; Yariv et al. 2020; Mildenhall et al. 2020; Schwarz et al. 2020; Niemeyer and Geiger 2021; Chan et al. 2021) require supervision based on multi-view data or synthesize low-resolution images, which limits the applicability of such methods. Other approaches (Gu et al. 2021; Chan et al. 2022; Or-El et al. 2022) explore 3D-consistent image generation under single-view training data, which are inspired by StyleGAN’s w latent code modulation

*These authors contributed equally.

†Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

network and design similar modulated neural radiance field network. These methods effectively solve the problem that training relies on multi-view data. But they model the whole image as an overall neural radiance field, which limits the semantic part manipulation of synthetic images. (Sun et al. 2022) introduces semantic segmentation maps into the rendering process, but it still uses a single rendering network and does not explicitly decouple the local semantic parts. It cannot directly edit the semantic regions of the synthetic image by manipulating the w latent code.

In this paper, we propose the first compositional neural radiance field for semantic 3D-aware portrait synthesis and manipulation, namely CNeRF. It divides the portrait image by semantic regions and learns an independent neural radiance field for each region. Therefore, we can independently control the local semantic regions of the synthetic 3D portrait by manipulating the latent code of the neural radiance field in each semantic region. Specifically, we split the overall generator network into k local semantic 3D generators. Each sub-network consists of modulated MLPs to output feature values, color values, mask values, and residual SDF representation of 3D space for each semantic region. Afterwards, the CNeRF conducts semantic part fusion and volume aggregation to render the complete 2D features, 2D colors and corresponding 2D masks. The 2D features and 2D masks can be further processed by a style-based 2D generator to obtain higher resolution portrait. The whole training process is divided into a low-resolution CNeRF rendering stage and a high-resolution image synthesis stage.

Our method is evaluated mainly on FFHQ dataset and corresponding semantic masks. While generating quality and diversity comparable to state-of-the-arts, our method can also finely manipulate semantic parts of generated portraits, and decouple the shape-texture of the semantic regions. In addition we evaluate our approach on a proprietary cartoon portrait dataset. High-quality generation and semantic part editing of the portraits demonstrate the application potential of our method.

To summarize, our main contributions are as following:

- We present the first compositional neural radiance field, which divides the image by semantic regions and learns an independent neural radiance field for each region.
- Based on CNeRF, we propose a semantic 3D-aware por-



Figure 1: The 3D-aware image synthesis and semantic manipulation. The left side shows the randomly generated face images, semantic masks, and 3D shapes. The right side shows the results of manipulating the local semantic regions in different views.

trait synthesis and manipulation method and achieve disentanglement of rendered shapes and textures within each semantic region.

- We propose the global and semantic discriminative loss and shape texture decoupling loss to address semantic entanglement during compositional rendering.
- Experiments reveal that our approach can finely manipulate the semantic parts, while maintaining high-fidelity 3D-consistent portrait synthesis.

Related Work

2D Generative Adversarial Networks

2D GANs (Goodfellow et al. 2014) have achieved remarkable success in computer vision. There are many approaches focused on exploring image generation for specific scenarios, including scene generation (Park et al. 2020; Arad Hudson and Zitnick 2021), face generation (Wiles, Koepke, and Zisserman 2018; Lee et al. 2020), human generation (Ma et al. 2021; Ren et al. 2022), etc. One of the most widely used generative frameworks is StyleGAN (Karras, Laine, and Aila 2019; Karras et al. 2020, 2021), whose w latent space provides a way to edit different attributes of the generated images. Some approaches explore the editing methods in StyleGAN’s latent space, such as InterFaceGAN (Shen et al. 2020), StyleFlow (Abdal et al. 2021), etc. Since StyleGAN is a noise-to-image type of generation method, it cannot input real images for editing. GAN inversion methods are proposed to solve the problem of real image editing using pre-trained StyleGAN, for instance, e4e (Tov et al. 2021), and PTI (Roich et al. 2021).

3D-Aware GANs

Recently, Neural Implicit Representations (NIR) are widely used for scene 3D modeling (Chabra et al. 2020; Chibane, Pons-Moll et al. 2020; Jiang et al. 2020), object shape and appearance rendering (Michalkiewicz et al. 2019; Niemeyer et al. 2019; Atzmon and Lipman 2020; Chibane, Alldieck, and Pons-Moll 2020; Gropp et al. 2020). Some methods (Sitzmann, Zollhöfer, and Wetzstein 2019; Mildenhall et al. 2020; Niemeyer et al. 2020) learn neural implicit representation by using multi-view 2D images without 3D data supervision. One of the most important methods is the Neural Radiance Fields (NeRF) (Mildenhall et al. 2020), which represents the 3D scene as a series of neural radiance and density fields and uses volume rendering (Kajiya and Von Herzen 1984) technique for 3D reconstruction. Many approaches recently focused on 3D-aware GAN under single-view image supervision. Pi-GAN (Chan et al. 2021) proposes a siren-based neural Radiance field and uses global latent code to control the generation of shapes and textures. GIRAFFE (Niemeyer and Geiger 2021) uses a two-stage rendering process, first generating a low-resolution feature map with a volume renderer, and then learning to boost the resolution of the output with a 2D CNN network. StyleNeRF (Gu et al. 2021) integrates NeRF into a style-based generator to improve rendering efficiency and 3D-consistency of high-resolution image generation. EG3D (Chan et al. 2022) proposes a tri-plane 3D representation method to improve rendering computational efficiency and generation quality. StyleSDF (Or-EI et al. 2022) merges a Signed Distance Fields 3D representation with a style-based 2D generator. Recently, some 3D-aware GAN methods (Sun et al. 2022;

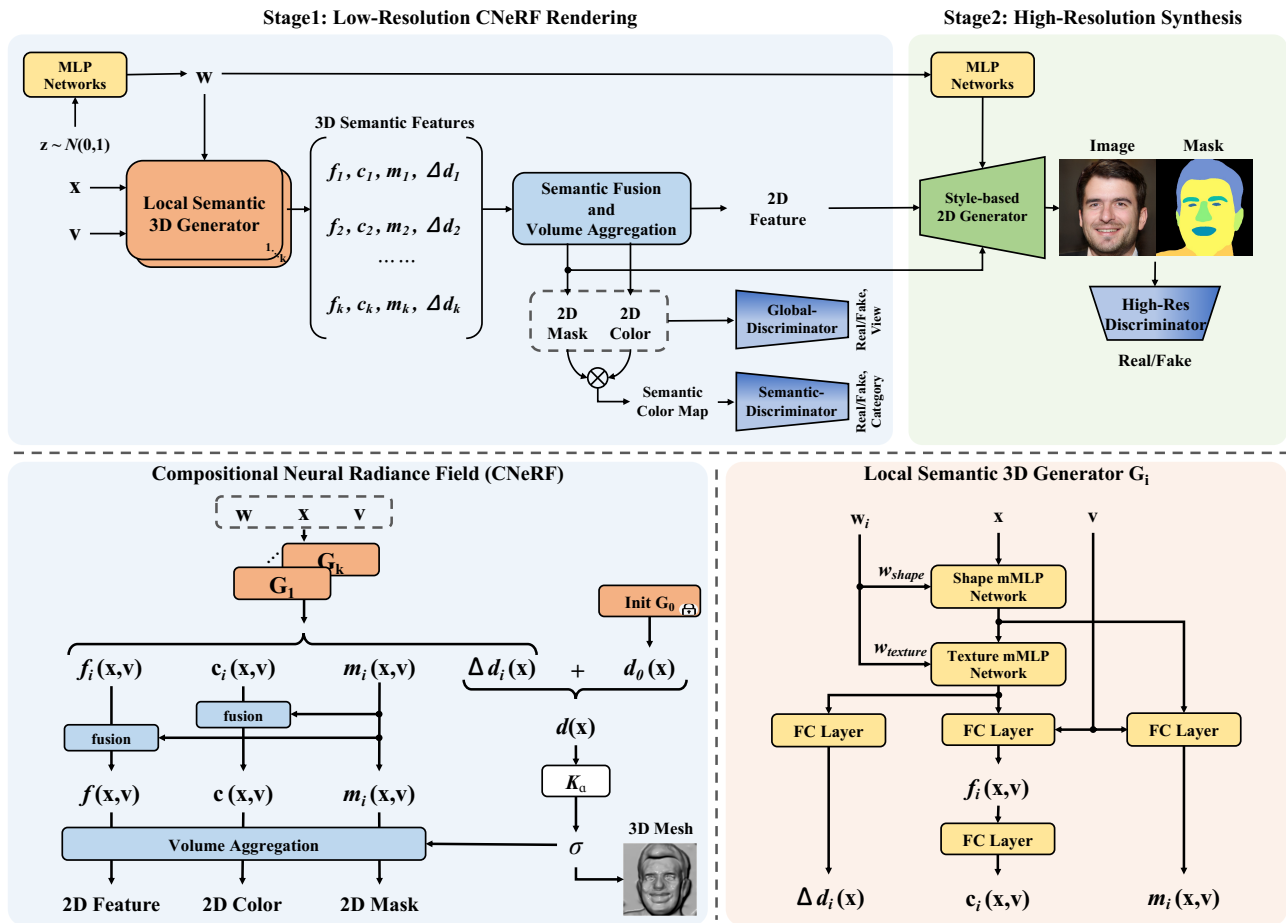


Figure 2: The overall framework of our approach. The training process is divided into a low-resolution CNeRF rendering stage and a high-resolution synthesis stage. The specific structure of CNeRF is shown at the bottom left. k local semantic 3D generators learn different semantic part neural radiance field and then sequentially performs semantic-based fusion and volume aggregation. The architectural details of each local semantic 3D generator is shown at the bottom right. The corresponding discriminators and loss functions are used in the two training stages, respectively.

Chen et al. 2022) introduce semantic segmentation into the generative network. But they render images and segmentation masks in single network with insufficient semantic part decoupling, which cannot directly edit the semantic regions. In this paper, our method explicitly decomposes the rendering network into multiple semantic region renderers for better semantic decoupling and semantic region editing.

Compositional Image Synthesis

Some methods (Arandjelović and Zisserman 2019; Azadi et al. 2020; Sbaji, Couprie, and Aubry 2021) generate complete scenes by combining different images as elements directly at the network input. Other approaches (Eslami et al. 2016; Yang et al. 2017; Burgess et al. 2019; Greff et al. 2019; Ehrhardt et al. 2020; Yang, Chen, and Soatto 2020; Arad Hudson and Zitnick 2021; Shi et al. 2022) explore the independent implicit representations of different objects in the scene image, which allows to manipulate the objects in the generated scene by controlling the latent code of each

object. For 3D-aware GAN methods, it is a novel idea to split the rendered image according to semantic regions and then design a compositional volume rendering method, which allows for individual manipulation of the semantic region in the generated 3D-consistent images.

Methodology

Overview

The overall model is divided into two stages. Firstly, we train a novel compositional neural radiance field to divide the image by semantic regions and learn an independent neural radiance field for each region, and finally fuses them and renders the complete 2D portrait. Secondly, we utilize a style-based 2D generator to upsample the low-resolution 2D images and obtain high-quality 3D-consistent portraits. In the inference, we can achieve semantic part control of the 3D-consistent portraits by manipulating the latent code of each semantic generator in CNeRF. The overall framework is shown in Figure 2, and the details are specified below.

Compositional Neural Radiance Field

We propose the Compositional Neural Radiance Field (CNeRF), as shown on the lower left of Figure 2. We split the overall scene renderer into k local semantic 3D generators, denoted as G_i ($i \in \{1, \dots, k\}$), with k representing the number of semantic categories we set according to the dataset. For the FFHQ dataset we set 9 local semantic categories, i.e., background, face, eyes, eyebrows, mouth, nose, ears, hair and the area below the head. Each local generator consists of modulated Multilayer Perceptron (mMLP) (Chan et al. 2021) and fully connected layers (FC Layers) of the same structure. The generator G_i takes as input the 3D coordinates \mathbf{x} , the view direction \mathbf{v} , and uses the latent code w_i learned from the noise of the standard Gaussian distribution as the network modulation parameters. G_i outputs two view-dependent color value $c_i(\mathbf{x}, \mathbf{v})$ and mask value $m_i(\mathbf{x}, \mathbf{v})$, where $c_i(\mathbf{x}, \mathbf{v})$ is obtained from the feature output $f_i(\mathbf{x}, \mathbf{v})$. In addition G_i outputs a view-independent residual SDF value $\Delta d_i(\mathbf{x})$, where SDF stands for Signed Distance Fields, a proxy of the volume density (Yariv et al. 2021; Or-EI et al. 2022). The network structure of G_i is shown in the lower right of Figure 2.

Local Semantic 3D Generator. For each local semantic 3D generator G_i , we further decouple the shape and texture properties by splitting the mMLP network into two parts to learn the shape and texture of each semantic part separately. To balance the model performance and the number of parameters, our shape network uses a 3-layer mMLP with SIREN activation (Sitzmann et al. 2020) and the texture network uses a 2-layer mMLP. The expression for j th layer of the mMLP is as follows,

$$\phi_j(x) = \sin(\gamma_j(W_j \cdot x + b_j) + \beta_j), \quad (1)$$

where W_j and b_j are network parameters, γ_j and β_j are modulation coefficients computed from the input w_i (w_{shape} or $w_{texture}$) latent code, and the mapping network for w_i is a 3-layer MLP with LeakyReLU activation. The w_i of each semantic generator may be the same or different at training iteration depending on random settings. Based on the above network, we use the FC layers to output different features at specific layers. Where in addition to the residual SDF value, other outputs are affected by the view direction.

Residual SDF Representation. We tried two types of volume density representation: direct volume density representation and SDF proxy. However, in our experiments we found that splitting the overall generator into local semantic generators destroys the learning process of object 3D shape and leads to incorrect 3D geometric reconstruction. We therefore propose a new volume density proxy, residual SDF representation, for compositional rendering. We add the $\Delta d_i(\mathbf{x})$ output from each semantic generator to a global SDF $d_0(\mathbf{x})$ initialized as a sphere to serve as the true SDF representation for the whole scene. This allows each G_i to learn only the changes of the SDF function value in corresponding region without learning the real physically meaningful SDF value, which makes good 3D geometric reconstruction possible. We transform the residual SDF values of the 3D semantic parts into the density field of the overall

scene following the equation below,

$$\begin{aligned} \sigma(\mathbf{x}) &= K_\alpha(d(\mathbf{x})) = \frac{1}{\alpha} \cdot \text{Sigmoid}\left(\frac{-d(\mathbf{x})}{\alpha}\right), \\ d(\mathbf{x}) &= d_0(\mathbf{x}) + \sum_i \Delta d_i(\mathbf{x}), \quad i \in \{1, \dots, k\}, \end{aligned} \quad (2)$$

where α is a learnable parameter to regulate the density tightness of the object surface (Or-EI et al. 2022).

Semantic-Based Fusion. In addition to learning a good global density representation, we need to fuse the color values output by G_i to form a complete color value. We utilize the idea of semantic mask-based fusion. G_i outputs the color c_i and feature f_i along with the mask m_i for the corresponding region, so we use m_i to perform a weighted fusion of c_i and f_i ,

$$\begin{aligned} c(\mathbf{x}, \mathbf{v}) &= \sum_i c_i(\mathbf{x}, \mathbf{v}) \cdot m_i(\mathbf{x}, \mathbf{v}), \quad i \in \{1, \dots, k\}, \\ f(\mathbf{x}, \mathbf{v}) &= \sum_i f_i(\mathbf{x}, \mathbf{v}) \cdot m_i(\mathbf{x}, \mathbf{v}), \quad i \in \{1, \dots, k\}, \end{aligned} \quad (3)$$

We do not limit the value range of m_i during the semantic fusion, but learn 2D masks end-to-end.

Volume Aggregation. For each pixel, we query the samples on a ray that originates at the camera position \mathbf{o} , and points at $\mathbf{r}(t) = \mathbf{o} + t\mathbf{v}$. We utilize the classical volume rendering equations to accumulate the samples on the ray,

$$\begin{aligned} \mathbf{C}(\mathbf{r}) &= \int_{t_n}^{t_f} T(t) \cdot \sigma(\mathbf{r}(t)) \cdot c(\mathbf{r}(t), \mathbf{v}) dt, \\ \mathbf{F}(\mathbf{r}) &= \int_{t_n}^{t_f} T(t) \cdot \sigma(\mathbf{r}(t)) \cdot f(\mathbf{r}(t), \mathbf{v}) dt, \\ \mathbf{M}_i(\mathbf{r}) &= \int_{t_n}^{t_f} T(t) \cdot \sigma(\mathbf{r}(t)) \cdot m_i(\mathbf{r}(t), \mathbf{v}) dt, \end{aligned} \quad (4)$$

where $T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right)$. We use the discrete form of the volume aggregation method following NeRF (Mildenhall et al. 2020), and obtain the rendered 2D color, 2D feature and k 2D semantic region masks.

Global and Semantic Discriminators

To help the training of our CNeRF model, we propose two discriminators, the Global-Discriminator (GD), which judges the overall rendering result, and the Semantic-Discriminator (SD), which judges the semantic parts. The GD takes as input the rendered 2D mask and 2D color image and outputs the predicted true/false label as well as the estimated view direction of the input generated image. The estimation of view direction is to ensure a good 3D-consistency of the synthesized image. The SD is designed to enhance semantic region disentanglement. In each iteration of training, the model randomly selects one of the k semantic categories and extracts the corresponding region in the 2D color image based on the generated 2D mask. The SD learns to discriminate the reality of the semantic part input and to predict the semantic category of the input image. Two discriminators consist of a residual CNN, with the network divided into two




Figure 3: Qualitative comparisons of 3D-consistent portrait synthesis between our method and state-of-the-art 3D-aware GAN methods on FFHQ dataset.

output layers at the end. The input of GD has two branches, image CNN and mask CNN. Then the output features of the two branches are added together. Besides, the input of SD is a single branch of semantic region image.

Loss Functions

Global and Semantic Discriminative Loss. For the adversarial losses, we utilize the non-saturating GAN loss with R1 regularization (Mescheder, Geiger, and Nowozin 2018) for our two discriminators, denoted as $\mathcal{L}_{adv}^{global}$ and $\mathcal{L}_{adv}^{semantic}$. In addition the two discriminators have additional view direction prediction and semantic category prediction, respectively. We use a smoothed $L1$ loss and a cross-entropy loss to constrain these two learning processes, denoted as $\mathcal{L}_{view}^{global}$ and $\mathcal{L}_{classify}^{semantic}$.

Shape Texture Decoupling Loss. Based on our local semantic 3D generator network, we propose a shape texture decoupling loss. It can help better decouple the shape-texture of each semantic region.

$$\begin{aligned} \mathcal{L}_{STD} = & \left\| \mathcal{M}(w_s^1, w_t^1) - \mathcal{M}(w_s^1, w_t^2) \right\|_1 \\ & + \max \left\{ \left\| \phi(I(w_s^1, w_t^1)) - \phi(I(w_s^2, w_t^1)) \right\|_1 \right. \\ & \left. - \left\| \phi(I(w_s^1, w_t^1)) - \phi(I(w_s^2, w_t^2)) \right\|_1 + m, 0 \right\}, \end{aligned} \quad (5)$$

where w_s and w_t are the abbreviation of w_{shape} and $w_{texture}$. The first half of the formula indicates the expectation that the generated 2D masks are consistent when we fix the w_{shape} and change the $w_{texture}$. The $\mathcal{M}(*, *)$ represents the rendered 2D mask. The second half represents a contrast loss, where we expect that inputting the same $w_{texture}$ will make the generated 2D color images closer in feature space than inputting two different $w_{texture}$. The ϕ is the middle feature layer of the pre-trained VGG network, and m is a fixed margin.

Eikonal Loss and Minimal Surface Loss. Eikonal Loss guarantees that the learned overall SDF of the object is physically valid (Gropp et al. 2020).

$$\mathcal{L}_{Eik} = \mathbb{E}_x (\| \nabla d(x) \|_2 - 1)^2. \quad (6)$$

Minimal Surface Loss is used to penalize the SDF value to ensure that the surface of the 3D object is smooth and free of false and invisible surfaces (Or-El et al. 2022).

$$\mathcal{L}_{MS} = \mathbb{E}_x (\exp(-100|d(x)|)). \quad (7)$$

The overall loss function is,

$$\begin{aligned} \mathcal{L}_{overall} = & \mathcal{L}_{adv}^{global} + \mathcal{L}_{view}^{global} + \mathcal{L}_{adv}^{semantic} + \mathcal{L}_{classify}^{semantic} \\ & + \mathcal{L}_{STD} + \mathcal{L}_{Eik} + \mathcal{L}_{MS}. \end{aligned} \quad (8)$$

The weights of each loss item are $\lambda_{view} = 15$, $\lambda_{Eik} = 0.1$, $\lambda_{MS} = 0.001$, and the rest are 1.

High-Resolution Synthesis

Direct synthesis of high-resolution 3D-consistent images is difficult, so we use a two-stage training approach (Gu et al. 2021; Or-El et al. 2022). First the 64×64 resolution 2D color, 2D feature map, and 2D masks are rendered. Then the CNeRF network weights are fixed and a style-based 2D generator is used to learn higher resolution image synthesis. We utilize the structure of StyleGAN2 (Karras et al. 2020) with input 2D feature, and add a mask input branch. We use the CNeRF's latent code w as input to the mapping network of the 2D generator. The 2D generator is shown in the upper right of Figure 2. We equip the 2D generator with a high-resolution discriminator, which has the same function as the Global-Discriminator of the previous stage.

Experiments

Datasets. We train and evaluate our model mainly on the FFHQ (Karras, Laine, and Aila 2019) dataset, which contains 70,000 high-quality portrait images with a maximum resolution of 1024×1024 . We use the Deeplabv3 (Chen et al. 2017) for semantic segmentation of these portrait images. In addition we utilize a 3D rendering pipeline to produce a high-quality cartoon portraits dataset, which contains 20,000 images at 1024 resolution. Similarly we perform facial semantic segmentation on this dataset. Our experimental re-

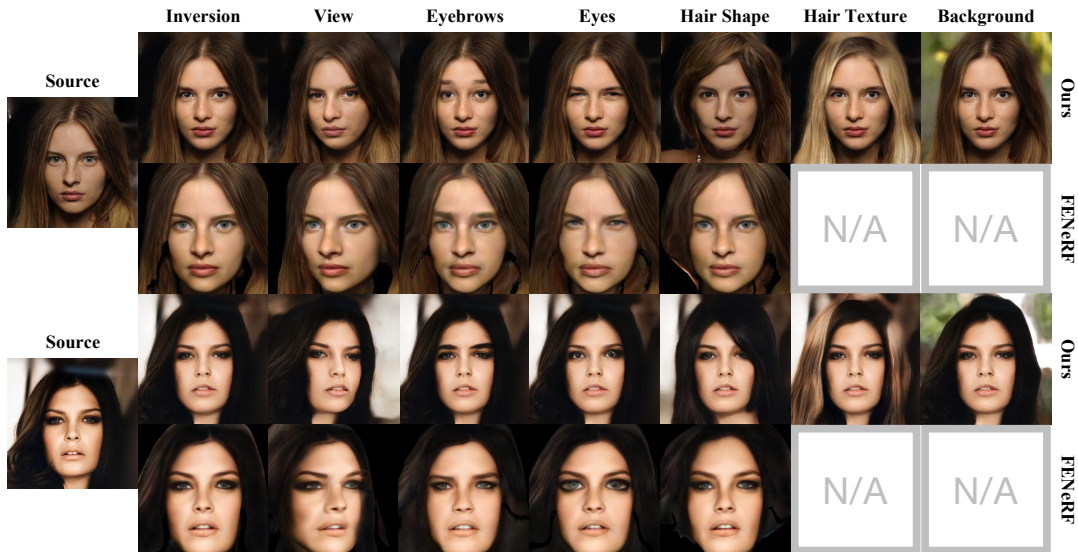


Figure 4: Comparisons of image inversion and semantic part manipulation for our method and FENeRF. Training on FFHQ dataset, inversion on CelebA-HQ (Liu et al. 2015) dataset. Perform GAN inversion using PTI (Roich et al. 2021).

	FID	KID	A	B	C
PiGAN (256)	83.0	85.8	✓		
StyleNeRF (256)	8.0	3.7	✓		
StyleSDF (256)	11.5	2.6	✓		
EG3D (512)	4.7	0.13	✓		
FENeRF (128)	28.2	17.3	✓	✓	
Ours (256)	13.2	1.2	✓	✓	✓
Ours (512)	12.3	0.21	✓	✓	✓

Table 1: Quantitative evaluation. All methods are evaluated under FFHQ dataset. The resolution of the synthesized image is shown in parentheses. A: 3D-consistent image synthesis, B: Semantic region manipulation, C: Semantic-level shape texture decoupling.

sults are mainly presented in these two datasets at 512 resolution.

Baselines. We compare with some of state-of-the-art 3D-aware GAN methods, including PiGAN (Chan et al. 2021), StyleNeRF (Gu et al. 2021), StyleSDF (Or-EI et al. 2022), EG3D (Chan et al. 2022), and FENeRF (Sun et al. 2022). Among them, FENeRF is also capable of semantic region editing of 3D-consistent faces, but it renders faces and segmentation masks in single network, which cannot directly edit the semantic regions.

Evaluation Metrics. We compute the two most commonly used image quality evaluation metrics, Frchet Inception Distance (FID) (Heusel et al. 2017) and Kernel Inception Distance (KID) (Bińkowski et al. 2018), which are used to evaluate the visual quality and diversity of the generated images. There is no unified and accurate evaluation metric for semantic manipulation effect, so we mainly perform qualitative comparison and demonstration.

Implementation Details. We train our model in two stages.

First, we train CNeRF using 64×64 resolution portrait images and segmentation masks, and render each output at 64×64 resolution. We use ADAM (Kingma and Ba 2014) optimizer with learning rates of 2×10^{-5} and 2×10^{-4} for the generator and discriminator respectively and $\beta_1 = 0, \beta_2 = 0.9$. Second, we freeze the CNeRF weights and train the 2D generator with same setup to StyleGAN2. Our experiments are carried out on 4 32GB Tesla V100 GPUs. Due to limited space more implementation details and detailed architecture of 2D generator and discriminators as well as more experimental results are presented in the supplementary material.

Comparisons

3D-Consistent Portrait Synthesis. Figure 3 shows the comparative results of 3D-consistent portrait generation. Our method is able to accurately reconstruct the 3D shape of the faces and also generate 2D masks corresponding to the faces. Our face 3D shape reconstruction is comparable to StyleSDF, but our method has semantic part manipulation ability that other comparison methods do not have. Table 1 presents the quantitative comparison results of our method with SOTA 3D-aware GAN methods, and the performance of our method is comparable to that of StyleSDF. The main contribution of our approach, however, is to synthesize 3D-consistent images while achieving good semantic-level property manipulation.

Semantic Region Manipulation. Figure 1 visually demonstrates the manipulation of the semantic parts of the synthesized faces, while ensuring 3D-consistent image synthesis. Our method can also control the shape and texture of the semantic regions separately, such as hair region in the figure. In addition figure 4 shows the comparison results of our method and FENeRF in real image semantic part manipulation. We use PTI (Roich et al. 2021) to train GAN inversion. Our method inverts more realistic face images

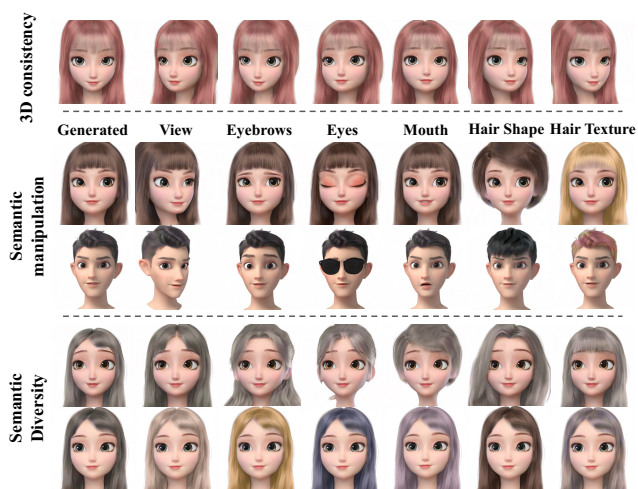


Figure 5: The 3D-consistent image synthesis and semantic region editing of cartoon portraits.

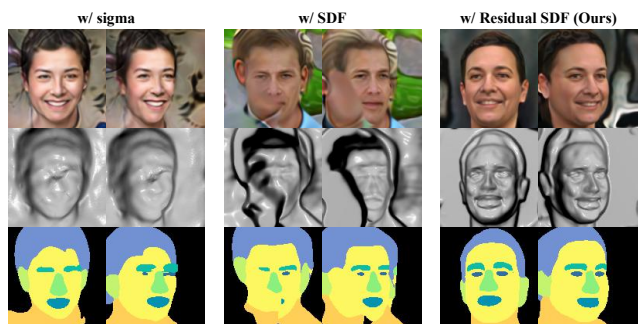


Figure 6: Ablation study of density function proxy representation in compositional rendering. Showing 128 resolution rendering results.

with better semantic part controllability. For example, our method can control the position as well as the shape of the eyebrows, while FENeRF can only edit at the original eyebrow area. Our model is able to decouple semantic region shape and texture properties, such as the displayed hair area, and change the background, which FENeRF cannot do. In addition our method manipulates the semantic parts by directly controlling the w_i latent code of each semantic generator, while FENeRF can only edit the mask manually and perform semantic editing with the help of GAN inversion.

Applications

Our method can be applied to the editing of real images, as shown in Figure 4. The quality of the image inversion and the semantic part manipulation effect illustrate the great potential of our method for application. In addition we test our method on a proprietary cartoon portrait dataset, as shown in Figure 5. High-quality 3D-consistent portrait synthesis as well as accurate and diverse semantic part editing demonstrate the potential of our approach to 3D animation, digital human and metaverse applications.




Figure 7: (a) Ablation of semantic discriminative loss. The second row represents the residual between the local manipulated image and the original image. (b) Ablation of shape texture decoupling loss. The last two rows show the manipulation of hair shape and texture.

Ablation Studies

We conduct ablation experiments on our CNeRF model from two aspects. For the overall 3D-consistent image synthesis ability, the residual SDF representation proposed in this paper plays a key role for compositional neural radiance field, as shown in Figure 6. When using the original volume density or the direct SDF proxy, our CNeRF produces incorrect 3D reconstruction and the synthesized images lose 3D-consistency. For the semantic part decoupling capability, we perform ablation experiments for the two correlated loss functions proposed in this paper, as shown in Figure 7. When using semantic discriminative loss (and Semantic-Discriminator), there is more accurate control over the semantic parts. When using shape texture decoupling loss, the model's performance is improved on the disentanglement of shape and texture properties and editing diversity.

Conclusion

In this paper, we propose a novel Compositional Neural Radiance Field for semantic 3D-aware portrait synthesis and manipulation. It enables high-quality 3D-consistent image synthesis while achieving independent manipulation of the semantic parts and decoupling the shape and texture within each semantic region. Qualitative and quantitative comparison experiments demonstrate the effectiveness and novelty of our method. The application performance of real image inversion and cartoon portrait 3D-aware editing show the potential of the method. There is still room to improve the quality of our approach in 3D reconstruction, and in future work we will explore the combination of our method with the latest 3D-aware GANs, such as EG3D.

Acknowledgments

This work is supported by the National Key Research and Development Program of China under Grant No. 2021YFC3320103.

References

- Abdal, R.; Zhu, P.; Mitra, N. J.; and Wonka, P. 2021. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (ToG)*, 40(3): 1–21.
- Arad Hudson, D.; and Zitnick, L. 2021. Compositional transformers for scene generation. *Advances in Neural Information Processing Systems*, 34: 9506–9520.
- Arandjelović, R.; and Zisserman, A. 2019. Object discovery with a copy-pasting gan. *arXiv preprint arXiv:1905.11369*.
- Atzmon, M.; and Lipman, Y. 2020. Sal: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2565–2574.
- Azadi, S.; Pathak, D.; Ebrahimi, S.; and Darrell, T. 2020. Compositional gan: Learning image-conditional binary composition. *International Journal of Computer Vision*, 128(10): 2570–2585.
- Bińkowski, M.; Sutherland, D. J.; Arbel, M.; and Gretton, A. 2018. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*.
- Burgess, C. P.; Matthey, L.; Watters, N.; Kabra, R.; Higgins, I.; Botvinick, M.; and Lerchner, A. 2019. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*.
- Chabra, R.; Lenssen, J. E.; Ilg, E.; Schmidt, T.; Straub, J.; Lovegrove, S.; and Newcombe, R. 2020. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *European Conference on Computer Vision*, 608–625. Springer.
- Chan, E. R.; Lin, C. Z.; Chan, M. A.; Nagano, K.; Pan, B.; De Mello, S.; Gallo, O.; Guibas, L. J.; Tremblay, J.; Khamis, S.; et al. 2022. Efficient geometry-aware 3D generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16123–16133.
- Chan, E. R.; Monteiro, M.; Kellnhofer, P.; Wu, J.; and Wetzstein, G. 2021. 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*, 5799–5809.
- Chen, L.-C.; Papandreou, G.; Schroff, F.; and Adam, H. 2017. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*.
- Chen, Y.; Wu, Q.; Zheng, C.; Cham, T.-J.; and Cai, J. 2022. Sem2NeRF: Converting Single-View Semantic Masks to Neural Radiance Fields. *arXiv preprint arXiv:2203.10821*.
- Chibane, J.; Alldieck, T.; and Pons-Moll, G. 2020. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6970–6981.
- Chibane, J.; Pons-Moll, G.; et al. 2020. Neural unsigned distance fields for implicit function learning. *Advances in Neural Information Processing Systems*, 33: 21638–21652.
- Ehrhardt, S.; Groth, O.; Monszpart, A.; Engelcke, M.; Posner, I.; Mitra, N.; and Vedaldi, A. 2020. RELATE: Physically plausible multi-object scene synthesis using structured latent spaces. *Advances in Neural Information Processing Systems*, 33: 11202–11213.
- Eslami, S.; Heess, N.; Weber, T.; Tassa, Y.; Szepesvari, D.; Hinton, G. E.; et al. 2016. Attend, infer, repeat: Fast scene understanding with generative models. *Advances in Neural Information Processing Systems*, 29.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Greff, K.; Kaufman, R. L.; Kabra, R.; Watters, N.; Burgess, C.; Zoran, D.; Matthey, L.; Botvinick, M.; and Lerchner, A. 2019. Multi-object representation learning with iterative variational inference. In *International Conference on Machine Learning*, 2424–2433. PMLR.
- Gropp, A.; Yariv, L.; Haim, N.; Atzmon, M.; and Lipman, Y. 2020. Implicit Geometric Regularization for Learning Shapes. In *International Conference on Machine Learning*, 3789–3799. PMLR.
- Gu, J.; Liu, L.; Wang, P.; and Theobalt, C. 2021. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Jiang, C.; Sud, A.; Makadia, A.; Huang, J.; Nießner, M.; Funkhouser, T.; et al. 2020. Local implicit grid representations for 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6001–6010.
- Kajiya, J. T.; and Von Herzen, B. P. 1984. Ray tracing volume densities. *ACM SIGGRAPH computer graphics*, 18(3): 165–174.
- Karras, T.; Aittala, M.; Laine, S.; Härkönen, E.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2021. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34: 852–863.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8110–8119.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lee, C.-H.; Liu, Z.; Wu, L.; and Luo, P. 2020. Maskgan: Towards diverse and interactive facial image manipulation.

- In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5549–5558.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep learning face attributes in the wild. *Proceedings of the IEEE international conference on computer vision*, 3730–3738.
- Ma, T.; Peng, B.; Wang, W.; and Dong, J. 2021. MUST-GAN: Multi-level statistics transfer for self-driven person image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13622–13631.
- Mescheder, L.; Geiger, A.; and Nowozin, S. 2018. Which training methods for GANs do actually converge? In *International conference on machine learning*, 3481–3490. PMLR.
- Michalkiewicz, M.; Pontes, J. K.; Jack, D.; Baktashmotlagh, M.; and Eriksson, A. 2019. Implicit surface representations as layers in neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4743–4752.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, 405–421. Springer.
- Niemeyer, M.; and Geiger, A. 2021. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11453–11464.
- Niemeyer, M.; Mescheder, L.; Oechsle, M.; and Geiger, A. 2019. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5379–5389.
- Niemeyer, M.; Mescheder, L.; Oechsle, M.; and Geiger, A. 2020. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3504–3515.
- Or-El, R.; Luo, X.; Shan, M.; Shechtman, E.; Park, J. J.; and Kemelmacher-Shlizerman, I. 2022. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13503–13513.
- Park, T.; Zhu, J.-Y.; Wang, O.; Lu, J.; Shechtman, E.; Efros, A.; and Zhang, R. 2020. Swapping autoencoder for deep image manipulation. *Advances in Neural Information Processing Systems*, 33: 7198–7211.
- Ren, Y.; Fan, X.; Li, G.; Liu, S.; and Li, T. H. 2022. Neural Texture Extraction and Distribution for Controllable Person Image Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13535–13544.
- Roich, D.; Mokady, R.; Bermano, A. H.; and Cohen-Or, D. 2021. Pivotal tuning for latent-based editing of real images. *arXiv preprint arXiv:2106.05744*.
- Sbai, O.; Couprie, C.; and Aubry, M. 2021. Surprising image compositions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3926–3930.
- Schwarz, K.; Liao, Y.; Niemeyer, M.; and Geiger, A. 2020. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems*, 33: 20154–20166.
- Shen, Y.; Gu, J.; Tang, X.; and Zhou, B. 2020. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9243–9252.
- Shi, Y.; Yang, X.; Wan, Y.; and Shen, X. 2022. Semantic-StyleGAN: Learning Compositional Generative Priors for Controllable Image Synthesis and Editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11254–11264.
- Sitzmann, V.; Martel, J.; Bergman, A.; Lindell, D.; and Wetzstein, G. 2020. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33: 7462–7473.
- Sitzmann, V.; Zollhöfer, M.; and Wetzstein, G. 2019. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems*, 32.
- Sun, J.; Wang, X.; Zhang, Y.; Li, X.; Zhang, Q.; Liu, Y.; and Wang, J. 2022. Fenerf: Face editing in neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7672–7682.
- Tov, O.; Alaluf, Y.; Nitzan, Y.; Patashnik, O.; and Cohen-Or, D. 2021. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4): 1–14.
- Wiles, O.; Koepke, A.; and Zisserman, A. 2018. X2face: A network for controlling face generation using images, audio, and pose codes. In *Proceedings of the European conference on computer vision (ECCV)*, 670–686.
- Yang, J.; Kannan, A.; Batra, D.; and Parikh, D. 2017. Lrgan: Layered recursive generative adversarial networks for image generation. *arXiv preprint arXiv:1703.01560*.
- Yang, Y.; Chen, Y.; and Soatto, S. 2020. Learning to manipulate individual objects in an image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6558–6567.
- Yariv, L.; Gu, J.; Kasten, Y.; and Lipman, Y. 2021. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34: 4805–4815.
- Yariv, L.; Kasten, Y.; Moran, D.; Galun, M.; Atzmon, M.; Ronen, B.; and Lipman, Y. 2020. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33: 2492–2502.