






Generative Deformable Radiance Fields for Disentangled Image Synthesis of Topology-Varying Objects

Ziyu Wang^{†1} , Yu Deng^{†2,3} , Jiaolong Yang³ , Jingyi Yu¹  and Xin Tong³ 

¹ ShanghaiTech University, ² Tsinghua University, ³ Microsoft Research Asia

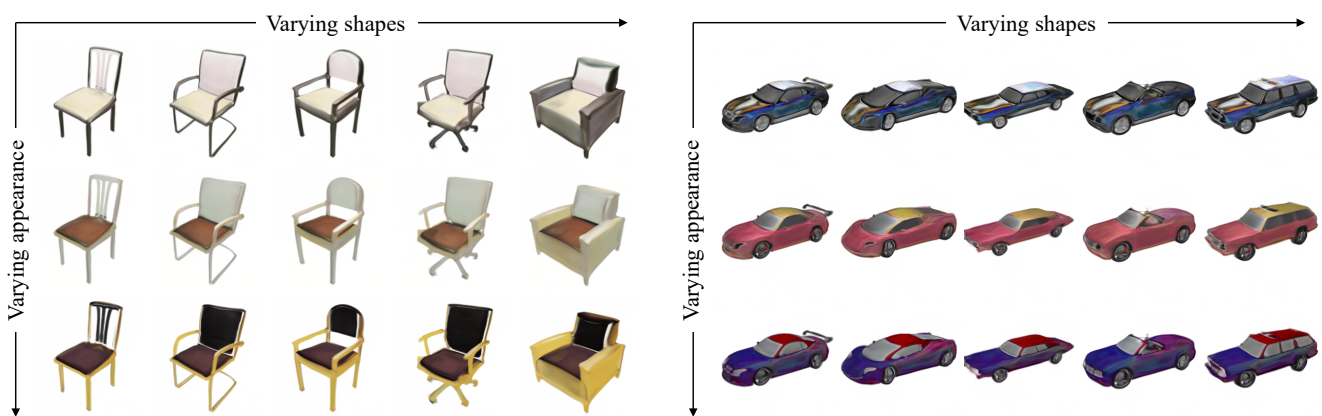


Figure 1: Image generated by our method with disentangled shape and appearance synthesis. Texture of the corresponding semantic components can be correctly transferred between different shapes with topology variations. See the accompanying video for multiview results.

Abstract

3D-aware generative models have demonstrated their superb performance to generate 3D neural radiance fields (NeRF) from a collection of monocular 2D images even for topology-varying object categories. However, these methods still lack the capability to separately control the shape and appearance of the objects in the generated radiance fields. In this paper, we propose a generative model for synthesizing radiance fields of topology-varying objects with disentangled shape and appearance variations. Our method generates deformable radiance fields, which builds the dense correspondence between the density fields of the objects and encodes their appearances in a shared template field. Our disentanglement is achieved in an unsupervised manner without introducing extra labels to previous 3D-aware GAN training. We also develop an effective image inversion scheme for reconstructing the radiance field of an object in a real monocular image and manipulating its shape and appearance. Experiments show that our method can successfully learn the generative model from unstructured monocular images and well disentangle the shape and appearance for objects (e.g., chairs) with large topological variance. The model trained on synthetic data can faithfully reconstruct the real object in a given single image and achieve high-quality texture and shape editing results.

CCS Concepts

• **Computing methodologies** → Rendering; Shape modeling; Image manipulation;

1. Introduction

3D-aware GANs [SLNG20, CMK*21, WCH*22] are capable of learning neural radiance field (NeRF) [MST*20] of an object category with topology variations, given only a collection of monoc-

[†] Work done when ZW and YD were interns at MSRA.

ular 2D images as supervision. With the advantages brought by the NeRF representation, they can easily generate images at arbitrary viewpoints and maintain multiview consistency when varying camera views, which greatly improves the prospects of GAN-based methods in applications like VR&AR.

However, a shortcoming of these methods is lack of controllability over factors beyond viewpoints, especially a disentangled control over shape and appearance. Such a limitation is not only seen in NeRF-based 3D-aware GANs, but also commonplace in existing methods that utilize NeRF representation for object modeling [LZZ*21, JA21]. This reduces their practical values when applied to image editing tasks such as texture transfer and shape manipulation. Although some methods [SLNG20, WCH*22] address this problem by designing separate latent codes for shape and appearance and send them to different layers of the NeRF generator, they still cannot guarantee independent changes because the predicted shape and color still share some common hidden features within the network.

Another line of works [BV99, LMR*15, GFK*18, JHTG20, ZYDL21, DYT21] for shape modeling have shown that deriving dense correspondence via deformations between object surfaces or their implicit fields can well address the disentanglement between appearance and shape. By learning reasonable correspondence between different shape surfaces, textures can be correctly transferred between corresponding semantic regions or structures, thus achieving high-quality disentanglement between shape and appearance control. However, most of these methods are trained with a 3D auto-decoding paradigm where ground truth 3D shapes with adequate geometry clues can be utilized for reasonable deformation learning. It is unclear how to model the dense correspondence between radiance fields and whether these methods can be adapted to the challenging 3D-aware GAN scenario without direct 3D supervision.

In this paper, we propose a novel generative deformable radiance fields that introduces 3D deformation into adversarial learning of NeRF-based GANs using only 2D images. With our method, disentangled control over shape and appearance can be achieved for image synthesis of objects with topology changes. The key idea is to encode the shapes of generated instances via deformation of a template radiance field shared across the category, and model the appearance within the shared template radiance field. Nevertheless, previous methods [GFK*18, ZYDL21, DYT21] have shown that pure deformation cannot handle correspondences for objects with topology difference and hence leads to undesired results for shape and texture editing. Therefore, we further add a correction field on top of the deformed radiance fields to handle structure changes and allow reasonable dense correspondence reasoning. To incorporate these modules into the 3D-aware GAN framework, we develop a set of carefully-designed loss functions to effectively learn the deformed radiance field with disentangled shape and appearance latent vector from unstructured monocular 2D images. With the learned generative deformable radiance fields, we further design an effective inversion initialization method to achieve high-quality image embedding and editing results.

We evaluate the proposed method on multiple datasets of topology-varying objects, including Photoshape [PRFS18],

ShapeNet [WSK*15], and CARLA [DRC*17, SLNG20]. We show that our method can not only generate 3D-consistent multiview images of these objects, but can also control their shapes and appearance in a disentangled manner, which greatly improves the controllability for NeRF-based image synthesis. In addition, we demonstrate the ability of our proposed image inversion scheme which can faithfully reconstruct real objects in an image and achieves high-fidelity texture editing and novel-view synthesis, even though we only trained our generative models on synthetic data. This reveals the strong potentials of our method for realistic virtual content creations and manipulations in future applications.

2. Related Work

Neural scene representations. Neural scene representation methods have witnessed significant progress in recent years, from the early black-box CNNs [KWKT15, DSTB16, TDB16, IZZE17, STH*19] to the latest neural implicit fields [PFS*19, MON*19, SZW19, SMB*20, MST*20, WLL*21] and hybrid representations [YFKT*22, PNM*20, CLC*22, PZX*21, LSS*21, MESK22]. A representative method among them is neural radiance field (NeRF) [MST*20], which leverages an MLP to predict color and density of each point in a 3D volumetric space for scene representation. NeRF and its follow-up works [BMT*21, BMV*22, PSB*21, TCY*22, SLNG20, CMK*21, LZZ*21, JA21] have demonstrated their strong capabilities in learning novel view synthesis of 3D scenes given only 2D images for training. However, they do not address explicit control over other properties of the scene such as shape and appearance. Several methods utilize two individual latent codes and send them to different layers of the radiance field network for certain degree of disentanglement between shape and appearance. They achieve color transfer at a global level but cannot handle detailed textures. Some recent methods adopt deformation fields for better shape and appearance disentanglement, but they either target at single scene reconstruction and editing [XXH*21] or object categories without large topology changes [XPMBB21, TPF*22]. We propose a new method to model topology-varying object categories and independently control the shape and appearance within a 3D-aware generative modelling scheme.

3D-aware image synthesis. Recent 3D-aware GANs [NPLT*19, SLNG20, CMK*21, NG21, DBS*21, DYXT22, CLC*22] are capable of learning multiview image generation of an object category from a monocular image set, which greatly simplifies the training configuration and demonstrates their potentials for large-scale image synthesis applications. Among them, images synthesized by NeRF-based generators [SLNG20, CMK*21, DBS*21] have demonstrated more strict 3D-consistency when varying camera views, which is an indispensable feature for the multiview image generation task. However, they still cannot achieve satisfactory control over shape and appearance. Perhaps the most relevant method to ours is in a very recent work of Tewari et al. [TPF*22] that is concurrent to ours. In this method, a deformation network is used to separate shape variations from a canonical radiance field that handles appearance changes, and an inverse deformation field is applied to enforce invertible mapping for reasonable deformation learning. However, this method does not deal with objects that

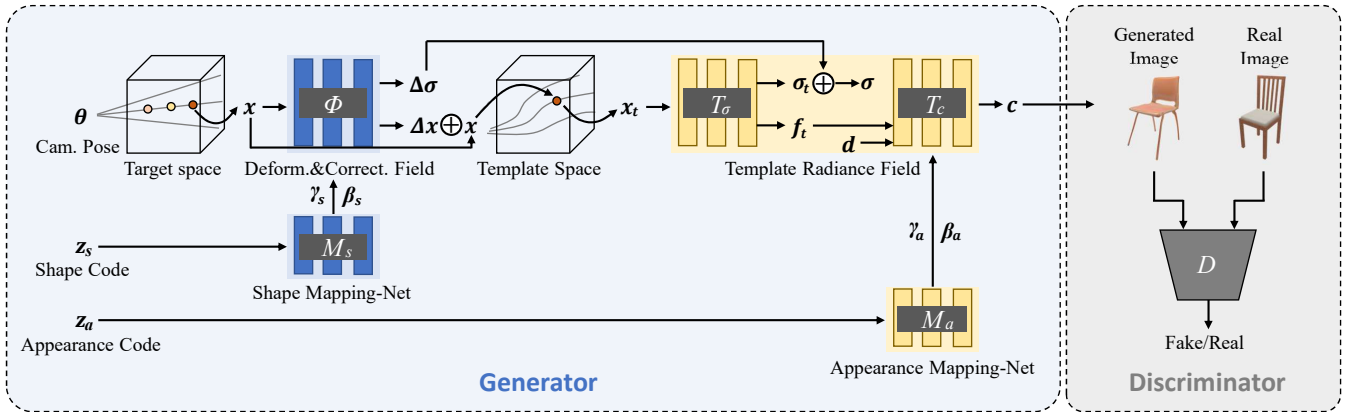


Figure 2: Overview of our generative deformable radiance field. It consists of a template radiance field conditioned on an appearance code z_a , and a deformation field and a correction field conditioned on a shape code z_s . A discriminator is introduced for adversarial learning between generated images and real ones. All networks are trained end-to-end on a collection of color images.

have structure variance. In this paper, we incorporate both deformation field and correction field into NeRF-based 3D-aware GAN to handle more complex objects such as chairs.

Shape modeling via deformations. Deformation-based modeling of a category of 3D shapes is a longstanding task in the literature and still an active topic to date. The deformations build dense correspondence between different shapes thus appearance of them can be easily defined on a shared template to achieve strong disentanglement with the geometry. Early works [BV99, VBPP06, ASK*05, PKA*09, LMR*15, ZKJB17] build linear subspace to characterize deformations between registered meshes. They are restricted to object categories with consistent topology and structures such as face and human body. Later works [GFK*18, YFST18] handle more complex objects by learning nonlinear mesh deformations from a large number of 3D shapes via deep neural network. Recently, some methods [DYT21, ZYDL21, ZYHC22, YTB*21] combine deformations with implicit shape representations. Among them, [DYT21] introduces an extra correction field to deal with structure discrepancies between shapes that are difficult to handle by continuous deformations. It shows promising results for detailed texture transfer between 3D shapes. Different from the aforementioned methods, we learn deformations and corrections in 3D space given only 2D images as training data, which is significantly more challenging than their 3D auto-decoding setup with ground truth 3D meshes for supervision.

Disentangled image synthesis and manipulations. A large volume of methods focus on disentangled control over multiple attributes of 2D image contents. Some methods [CDH*16, HMP*16] propose to learn disentangled representation by 2D CNNs in a fully unsupervised manner but often cannot guarantee disentanglement between semantically meaningful attributes. More of existing works [DSTB16, TYL17, PAM*18, SLY*18, PLWZ19, DYC*20, AZMW21] leverage certain kinds of attribute labels as guidance to achieve controllable image synthesis and editing of desired properties. Nevertheless, they show inferior results for 3D pose disentanglement compared to methods based on state-of-the-art 3D representations. Our method utilizes the NeRF representation thus can

not only achieve disentangled image manipulation of shape and appearance but also synthesize their novel views with high 3D consistency.

3. Overview

Given a monocular image collection of an object category with topology variations, we aim to learn an image generator G that can synthesize multiview images of virtual instances and control their shape and appearance attributes independently. Specifically, given a shape code $z_s \in \mathbb{R}^{d_s}$, an appearance code $z_a \in \mathbb{R}^{d_a}$, and a camera pose $\theta \in \mathbb{R}^3$ sampled from prior distributions as input, the generator G outputs an image I bearing the corresponding properties:

$$G: (z_s, z_a, \theta) \in \mathbb{R}^{d_s+d_a+3} \rightarrow I \in \mathbb{R}^{H \times W \times 3}. \quad (1)$$

To achieve 3D-consistent image generation given different camera poses, we adopt radiance field as our underlying 3D representation and follow a typical 3D-aware GAN training paradigm as in [CMK*21]. To achieve separate control of shape and appearance during image synthesis, we propose a novel generative deformable radiance field to represent these two attributes in a disentangled manner, as shown in Fig. 2. The generative deformable radiance field consists of a template radiance field conditioned on the appearance code z_a that handles texture variations (Sec. 4.1), together with a deformation field and a correction field dealing with shape variations (Sec. 4.2), both conditioned on the shape code z_s . With the above architecture, an image can be obtained via a standard volumetric rendering process (Sec. 4.3). Considering that the whole framework is trained with a weak image-level adversarial loss, we further introduce a set of dedicated loss functions as the regularization of the two fields to ensure the deformation field and correction field characterize reasonable shape changes (Sec. 4.4). Based on our learned generator, we develop a novel image inversion scheme that faithfully recovers the input image meanwhile maintains the inherent properties of the underlying deformable radiance field representation thus capable of high quality shape and texture editing (Sec. 5).

4. Generative Deformable Radiance Field

Our generative deformable radiance field models a category of topology-varying objects and synthesizes images of new instances with a disentangled control over shape, appearance, and camera viewpoints. It consists of a template radiance field, a deformation field, and a correction field, described as follows.

4.1. Template Radiance Field

The template radiance field T determines the appearance of objects in a shared template volumetric space. It is represented by an MLP which takes an appearance code \mathbf{z}_a as input, along with a 3D position $\mathbf{x}_t \in \mathbb{R}^3$ in the template space and a view direction $\mathbf{d} \in \mathbb{R}^3$, and predicts the color and density $(\mathbf{c}, \sigma_t) \in \mathbb{R}^4$ of \mathbf{x}_t :

$$T : (\mathbf{z}_a, \mathbf{x}_t, \mathbf{d}) \in \mathbb{R}^{d_a+6} \rightarrow (\mathbf{c}, \sigma_t) \in \mathbb{R}^4. \quad (2)$$

The network structure of T is shown in Fig. 2. The first several layers of the network are not conditioned on the appearance code, which are used to generate a shared density field for all objects within the category as their template shape. To characterize appearance changes of the template, we inject the appearance code into the FiLM-SIREN blocks [CMK*21] starting from an intermediate layer using modulations of frequencies γ_a and phase shifts β_a produced by an appearance code mapping network M_a (see [CMK*21] for details). With the above network structure, Eq. (2) can be re-written as:

$$T_\sigma : \mathbf{x}_t \in \mathbb{R}^3 \rightarrow (\mathbf{f}_t, \sigma_t) \in \mathbb{R}^{d_f+1}. \quad (3)$$

$$T_c : (\mathbf{z}_a, \mathbf{f}_t, \mathbf{d}) \in \mathbb{R}^{d_a+d_f+3} \rightarrow \mathbf{c} \in \mathbb{R}^3, \quad (4)$$

where \mathbf{f}_t is the intermediate feature, and T_σ and T_c denote the two sub-nets.

4.2. Deformation Field and Correction Field

The deformation field and the correction field are responsible for modeling shape variations within the class. As our underlying representation is an implicit one, we define an “inverse” deformation field, which deforms a 3D position in the target space (i.e., the object space for image rendering) into the shared template space. In addition, since both the deformation and correction are used to derive shape (density) changes given the shape code \mathbf{z}_s , we represent them via a single network Φ . The network takes the shape code \mathbf{z}_s and a 3D point $\mathbf{x} \in \mathbb{R}^3$ in the target space as input, and predicts a corresponding deformation vector $\Delta\mathbf{x} \in \mathbb{R}^3$ and a density correction $\Delta\sigma \in \mathbb{R}$:

$$\Phi : (\mathbf{z}_s, \mathbf{x}) \in \mathbb{R}^{d_s+3} \rightarrow (\Delta\mathbf{x}, \Delta\sigma) \in \mathbb{R}^4. \quad (5)$$

Φ is also an MLP with a set of FiLM-SIREN blocks [CMK*21] where the hidden layer features are modulated by frequencies γ_s and phase shifts β_s produced by a shape code mapping network M_s . With the predicted deformation $\Delta\mathbf{x}$ and correction $\Delta\sigma$, the final density $\sigma \in \mathbb{R}$ and color $\mathbf{c} \in \mathbb{R}^3$ of the target space point \mathbf{x} with view direction \mathbf{d} can be obtained via:

$$\sigma(\mathbf{z}_s, \mathbf{x}) = T_\sigma^{[\sigma_t]}(\mathbf{x} + \Phi^{[\Delta\mathbf{x}]}(\mathbf{z}_s, \mathbf{x})) + \Phi^{[\Delta\sigma]}(\mathbf{z}_s, \mathbf{x}), \quad (6)$$

$$\mathbf{c}(\mathbf{z}_s, \mathbf{z}_a, \mathbf{x}, \mathbf{d}) = T_c(\mathbf{z}_a, T_\sigma^{[\mathbf{f}_t]}(\mathbf{x} + \Phi^{[\Delta\mathbf{x}]}(\mathbf{z}_s, \mathbf{x})), \mathbf{d}), \quad (7)$$

where the superscripts denote the corresponding outputs of the networks when there are multiple ones. In a nutshell, the deformation field deforms a point in the target space to the template radiance field to obtain its density and color, and meanwhile the correction field modifies its density to handle topology variations that cannot be well modeled by deformation. It is worth mentioning that [DYT21] also combines deformation and correction, but their method only models 3D shape without considering appearance, and it uses an auto-decoder framework to learn signed distance fields for a given 3D mesh collection. In contrast, we learn 3D deformation and correction using unstructured 2D images with generative radiance field modeling and we handle both shape and appearance.

It can be seen from Eq. 6 that our deformable radiance field architecture *completely disentangles shape variations from appearance*, as the final density of a point only depends on the shape code \mathbf{z}_s and is not affected by \mathbf{z}_a . Also note that although the final color of the point \mathbf{x} appears to also depend on the shape code \mathbf{z}_s per Eq. (7), it is actually determined by the color code \mathbf{z}_a and the deformed position of \mathbf{x} in the template space. As a result, as long as Φ can deform the points on different shapes that are correspondences to the same position in the template space, a desired *full disentanglement between shape and appearance control can be achieved*. In Sec. 4.4, we introduce several regularizations to learn such a deformation.

4.3. Image Rendering

We render images of the radiance field via a standard volumetric rendering process as in [MST*20, CMK*21]. Given a camera pose θ , we cast rays in the target object space following the camera intrinsics and extrinsics and sample points along each ray using the hierarchical sampling strategy of [MST*20]. For a camera ray \mathbf{r} with sampled points $\{\mathbf{x}_i\}$ sorted from near to far, its corresponding color $C(\mathbf{r})$ can be obtained via:

$$C(\mathbf{r}) = \sum_{i=1}^n W(\mathbf{x}_i)(1 - \exp(-\sigma(\mathbf{x}_i)\delta_i))\mathbf{c}(\mathbf{x}_i, \mathbf{d}_i), \quad (8)$$

$$W(\mathbf{x}_i) = \exp\left(-\sum_{j=1}^{i-1} \sigma(\mathbf{x}_j)\delta_j\right), \quad (9)$$

where $\sigma(\mathbf{x}_i)$ and $\mathbf{c}(\mathbf{x}_i, \mathbf{d}_i)$ are corresponding density and color of \mathbf{x}_i obtained via Eq. (6) and (7), and δ_i is the distance between \mathbf{x}_i and its adjacent point. Note the shape and color codes are omitted here for brevity.

4.4. Training Strategy

We follow a typical 3D-aware GAN’s training strategy to learn our generative deformable radiance field using unstructured 2D image collection. Specifically, we sample shape code \mathbf{z}_s , color code \mathbf{z}_a , and camera pose θ from prior distributions and generate corresponding images $G(\mathbf{z}_s, \mathbf{z}_a, \theta)$ via Eq. (6)–(9). We send the generated images as well as real images sampled from the training data to a discriminator D and apply a non-saturating GAN loss with R1 regularization [KLA19, MGN18] and a pose regularization following [DYXT22]. To ensure that the deformation field derives reasonable correspondences for shape and appearance disentanglement, we further design several regularizations described below.




Figure 3: Illustration of the normal consistency regularization.

Normal consistency. Inspired by [DYT21], we enforce a normal consistency regularization between the target space radiance field and the template radiance field to ensure that the corresponding points in the two spaces derived by the deformation field have similar normal directions. For objects in a same category, their corresponding surface points with similar semantic meaning usually have similar surface normals. Therefore, the normal consistency regularization encourages the deformation field to establish correspondence between points sharing similar semantic meaning. However, our radiance fields do not characterize well-defined object surfaces compared to the signed distance field in [DYT21]. To tackle this problem, we propose a new normal consistency regularization tailored for our radiance field representation (Fig. 3):

$$L_{\text{normal}} = \sum_{r=1}^R \sum_{\mathbf{x}_i \in \mathcal{N}(\tilde{\mathbf{x}}^r)} (1 - \langle \text{sg}(\nabla \sigma(\mathbf{x}_i)), \nabla T_{\sigma}^{[\sigma]}(\mathbf{x}_i + \Delta \mathbf{x}_i) \rangle), \quad (10)$$

where R is the total number of rays. $\Delta \mathbf{x}_i$ is the deformation of point \mathbf{x}_i as described in Eq. (5) which derives the correspondence between \mathbf{x}_i and $\mathbf{x}_i + \Delta \mathbf{x}_i$. For σ and $\Delta \mathbf{x}_i$, the dependency on latent code is omitted for brevity. ∇ denotes the spatial gradient of a volumetric field, which can be effectively calculated via auto-gradient of a deep learning framework such as Pytorch [PGM*19], sg denotes the *stop-gradient* operator, $\langle \cdot, \cdot \rangle$ is the cosine distance, and $\mathcal{N}(\tilde{\mathbf{x}}^r)$ denotes the neighborhood of the approximated surface point $\tilde{\mathbf{x}}^r$ along ray \mathbf{r} . The approximated surface point can be obtained via weighted average over all points $\{\mathbf{x}_i\}$ along the ray [MST*20]:

$$\tilde{\mathbf{x}}^r = \sum_{i=1}^n W(\mathbf{x}_i) (1 - \exp(-\sigma(\mathbf{x}_i) \delta_i)) \mathbf{x}_i. \quad (11)$$

We only apply the normal consistency loss for points near the approximated surface as the normals in the free space of a radiance field can be arbitrary. In practice, We adopt a hierarchical sampling strategy following [MST*20] and only select 5% of total samples that are closest to the surface point to calculate this loss. Moreover, we do not propagate the gradient back to the final density $\sigma(\mathbf{x}_i)$ to ensure that the loss only regulates the learned deformation field and template radiance field.

Deformation smoothness. We also apply a smoothness constrain on the deformation field to avoid irregular deformation:

$$L_{\text{smooth}} = \sum_{i=1}^n \|\nabla \Phi^{[\Delta \mathbf{x}]}(\mathbf{x}_i)\|_2. \quad (12)$$

where $\|\cdot\|_2$ denotes the L_2 norm.

Deformation rigidity. We further apply a new rigidity regulariza-

tion to ensure that the learned deformation is locally rigid:

$$L_{\text{rigid}} = \sum_{i=1}^n \|J_i^T J_i - I\|_F, \quad (13)$$

where $J_i = \nabla(\mathbf{x}_i + \Phi^{[\Delta \mathbf{x}]}(\mathbf{x}_i))$ is the Jacobian matrix of the deformed point in the template space with respect to the original point in the target space. $\|\cdot\|_F$ denotes the Frobenius norm.

Minimal correction. We follow [DYT21] to impose a minimal correction prior to encourage shape modeling via deformation rather than correction:

$$L_{\text{correction}} = \sum_{i=1}^n |\Delta \sigma(\mathbf{x}_i)|. \quad (14)$$

5. Real Image Inversion and Editing

Although our generative deformable radiance field is learned in a GAN training paradigm, it can be applied to disentangled editing of real images. To this end, we propose a novel image inversion scheme that can faithfully reconstruct the details of a given image meanwhile maintaining the shape and appearance disentanglement for image editing task. Figure 4 shows the overview of the method.

5.1. Inversion Strategy

Given a real image \hat{I} of an object with camera pose θ , we aim to find its corresponding deformable radiance field representation that can not only recovery the image content under the given viewpoint, but also support disentangled editing of shape and appearance and produce plausible results under novel views. To achieve this, a straightforward solution is to find the shape code \mathbf{z}_s and color code \mathbf{z}_a that describe the real image. However, obtaining such codes with high-fidelity reconstruction result for a real image is almost impossible as shown by previous GAN inversion methods [AQW19, KLA*20, Zho21]. On the other hand, representing the image with intermediate features produces more faithful reconstruction result but may sacrifice the image editing quality [ZSZZ20, CMK*21].

Our generative deformable radiance field naturally disentangles shape and appearance for image editing. Therefore, we only need to guarantee that the recovered shape is reasonable for novel view synthesis and the recovered appearance is of high fidelity. We achieve this by optimizing the following objective function:

$$\begin{aligned} \operatorname{argmin}_{\mathbf{w}_s, \boldsymbol{\gamma}_a, \boldsymbol{\beta}_a} \|I(\mathbf{w}_s, \boldsymbol{\gamma}_a, \boldsymbol{\beta}_a, \theta) - \hat{I}\|_2^2 + \text{LPIPS}(I(\mathbf{w}_s, \boldsymbol{\gamma}_a, \boldsymbol{\beta}_a, \theta), \hat{I}) \\ + \|\mathbf{w}_s - \mathbf{w}_s^0\|_2^2 + \|\boldsymbol{\gamma}_a - \boldsymbol{\gamma}_a^0\|_2^2 + \|\boldsymbol{\beta}_a - \boldsymbol{\beta}_a^0\|_2^2, \end{aligned} \quad (15)$$

where \mathbf{w}_s is the last hidden feature of the shape mapping network, $\boldsymbol{\gamma}_a$ and $\boldsymbol{\beta}_a$ are output features of the appearance mapping network which are also the frequencies and phase shifts of each FILM-SIREN block of the template field network, and \cdot^0 denotes the initial features for optimization (see Sec. 5.2). $I(\mathbf{w}_s, \boldsymbol{\gamma}_a, \boldsymbol{\beta}_a, \theta)$ denotes the generated image and $\text{LPIPS}(\cdot, \cdot)$ is the perceptual loss from [ZIE*18]. We optimize the hidden features of the mapping network instead of the frequencies and phase shifts for shape to ensure that the recovered 3D geometry is more plausible under novel

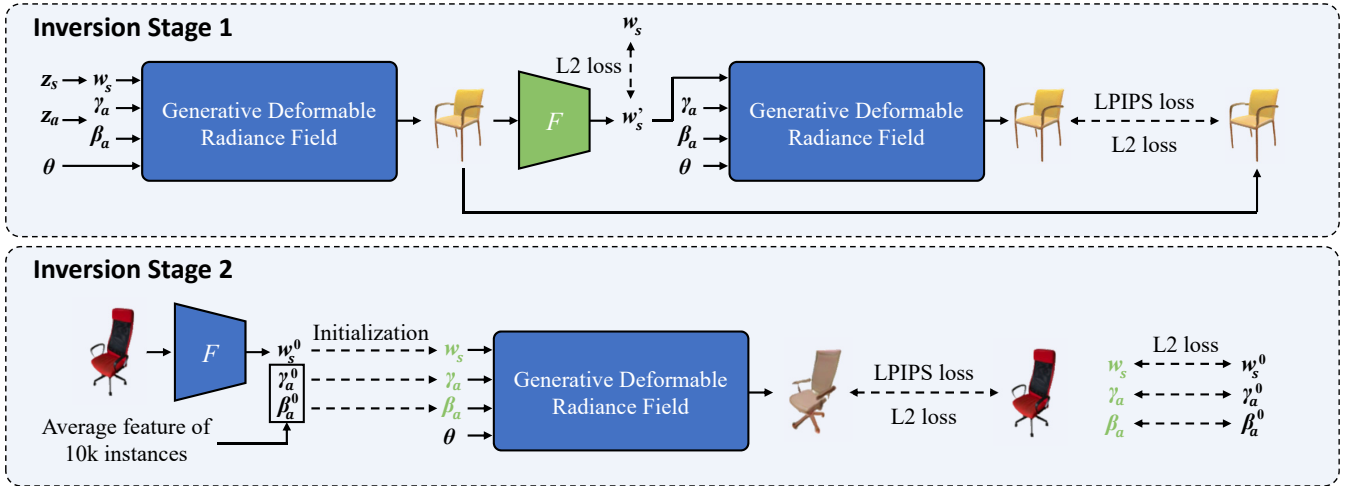


Figure 4: Overview of our image inversion pipeline. In the first stage, we train a 2D CNN using pairs of generated images and their corresponding shape and appearance features, where w_s denotes the last hidden feature of the shape mapping network. In the second stage, given a real image, we use the pre-trained 2D CNN to predict a shape feature as initialization. Then, we optimize the shape and appearance features to faithfully reconstruct the given image. Green color indicates the components to be optimized in each stage.

views, and we directly optimize the frequencies and phase shifts for appearance to obtain high-fidelity textures.

5.2. Inversion Initialization

For object categories with diverse structures such as chairs, a proper initialization of the features to be solved in the inversion optimization is important. The following feature initialization strategy is proposed in this work. First, we randomly generate a collection of images $\{I^i\}$ using our trained deformable radiance field generator and record their corresponding camera poses $\{\theta^i\}$ and features $\{w_s^i\}$, $\{\gamma_a^i\}$, and $\{\beta_a^i\}$. Then, we train a CNN F that predicts feature w_s using the following loss function:

$$\begin{aligned} \operatorname{argmin}_F & \|F(I^i) - w_s^i\|^2 + \|I(F(I^i), \gamma_a^i, \beta_a^i, \theta^i) - I^i\|^2 \\ & + \text{LPIPS}(I(F(I^i), \gamma_a^i, \beta_a^i, \theta^i), I^i), \end{aligned} \quad (16)$$

which enforces F to reproduce the shape features w_s^i and the image I^i . After training, given a real image \hat{I} , we predict its corresponding shape feature $F(\hat{I})$ as the initial value w_s^0 in Eq. (15). For the appearance features, we empirically found that their initialization has minor influence on the inversion result, so we directly sample 10K random appearance codes z_a and calculate their average frequencies and phase shifts as the initial values γ_a^0 and β_a^0 in Eq. (15).

5.3. Disentangled Image Editing

Given a real image \hat{I} , we can first obtain its corresponding shape and appearance features w_s , γ_a , and β_a via Eq. (15). Then, we can achieve a disentangled editing over its shape and appearance by simply replacing the features with those of another image, and synthesize multiview images of the editing results given arbitrary camera viewpoints. For instance, given two real images \hat{I}^i and \hat{I}^j , we can easily obtain a new editing result $I(w_s^i, \gamma_a^j, \beta_a^j, \theta)$, which has the shape of \hat{I}^i and the appearance of \hat{I}^j .

6. Experiments

6.1. Implementation Details

We evaluate our methods on four datasets: PhotoShape chair [PRFS18], CARLA [SLNG20, DRC*17], ShapeNet car, and ShapeNet plane [CFG*15]. For PhotoShape chair, we render a single view image for each chair instance with a random camera viewpoint sampled from the upper hemisphere with 30 degrees pitch variation, and obtain 15K images in total. For CARLA, we use the 10K rendered images provided by the dataset. For ShapeNet car and plane, we render three views for each object and obtain 10K and 12K total images respectively, following a similar camera distribution as in PhotoShape chair. Our deformation network is an MLP with 4 layers with the FiLM-SIREN structure [CMK*21]. The template network has 5 layers with sine activation for density generation, and another 4 FiLM-SIREN layers for color output. Our discriminator is adapted from [CMK*21], and the inversion initialization CNN is a ResNet-34 network. During training, we sample shape code z_s and appearance code z_a from a normal distribution of dimension 256. We use a single random code for both shape and appearance of a generated instance. We randomly sample camera pose θ from the same distribution used to generate training data. We use Adam optimizer and train the model progressively from 64×64 resolution to 128×128 with batchsize 32 on 8 Nvidia Tesla V100 GPUs with 32GB memory. The training takes 3 to 6 days.

6.2. Random Generation Results

Figure 5 shows the random image generation results by our method on four different datasets. Our method can generate images of various objects with diverse structures and freely change camera viewpoints. The multiview generation results can be found in the accompanying video.



Figure 5: Random image generation results of our method on four different datasets. (Best viewed with zoom-in)

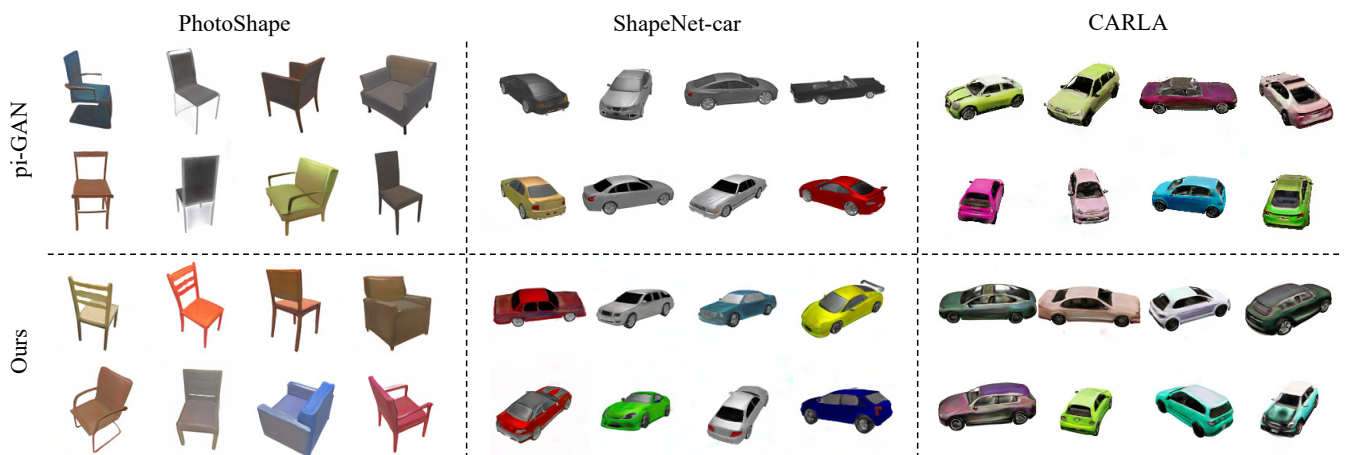


Figure 6: Qualitative comparison with the 3D-aware GAN model *pi*-GAN. Note that *pi*-GAN does NOT address shape-appearance disentanglement as we do, and that our method is trained on the same data without extra labels for disentanglement.

6.3. Disentangled Generation Results

Our method enables a disentangled control over shape and appearance attributes of the generated images. We achieve this by varying either the shape code or appearance code and keeping the other unchanged. The corresponding generation results are shown in Fig. 1 and Fig. 7 (see video for multiview results). With our generative deformable radiance field and the carefully designed training strategy, we obtain plausible dense correspondence between different objects. Textures can be correctly transferred between shape structures of similar semantic meaning. Note that our method is only trained with 2D images without any semantic labels.

In Fig. 8, we illustrate the the learned template fields for dif-

Table 1: Quantitative comparison with *pi*-GAN.

Methods	PhotoShape 128		ShapeNet-car 128		CARLA 128	
	FID	KID	FID	KID	FID	KID
<i>pi</i> -GAN	19.4	0.71	18.9	0.81	41.5	2.10
Ours	23.8	1.09	27.7	1.54	41.2	2.07

ferent categories and how different shapes are generated through deformation and correction.

We further show the latent space interpolation results in Fig. 9. Our disentanglement design makes it possible to interpolate either shape and appearance while keeping the other unchanged. The in-




Figure 7: Qualitative comparison of code swapping results on CARLA. The images on the diagonal are the reference instances and others are generated by combining the row-wise shape code and column-wise appearance code. Our method achieves much better shape and appearance disentanglement than EditNeRF and CodeNeRF (e.g., see the color of car roof in the column-wise texture swapping results)

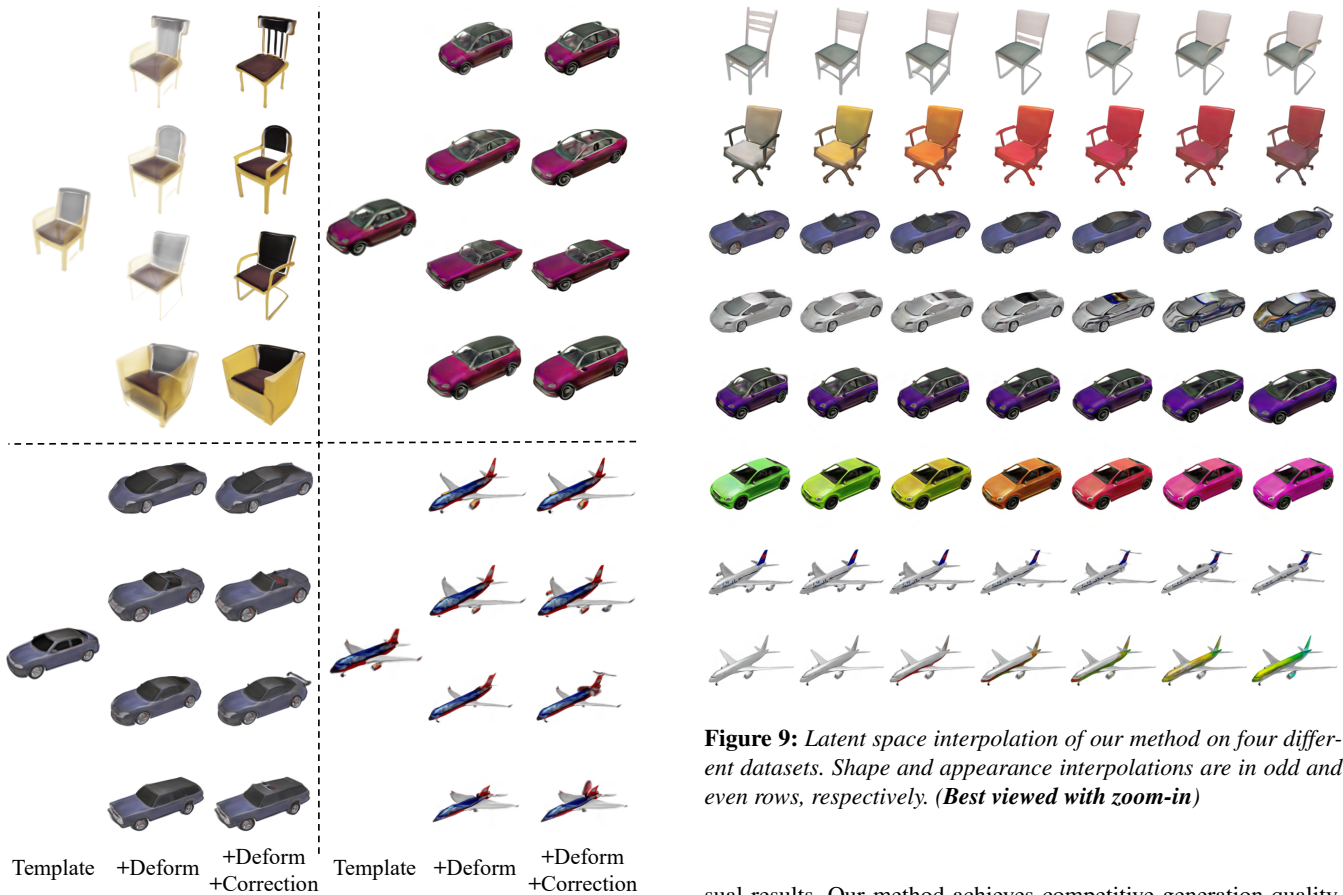


Figure 8: The learned templates and generated shapes by our method. (Best viewed with zoom-in)

terpolation produces reasonable intermediate results for two shapes with topology difference.

6.4. Comparison with Prior Art

Comparison with pi-GAN. First, we compare our uncured generation results with pi-GAN [CMK*21], and Fig. 6 shows the vi-

Figure 9: Latent space interpolation of our method on four different datasets. Shape and appearance interpolations are in odd and even rows, respectively. (Best viewed with zoom-in)

sual results. Our method achieves competitive generation quality, but further enables shape and appearance disentanglement which cannot be achieved by pi-GAN. Table 1 further shows the quantitative comparison, where we evaluate the FID [HRU*17] and KID [BSAG18] scores using 5K real images and 5K generated images. Compared to pi-GAN, our method achieves disentangled control of shape and appearance with only a slight increase of FID and KID.

Comparison with CodeNeRF and EditNeRF. To further validate the disentangled image synthesis ability of our method and compare with the state-of-the-art methods, we conduct shape and

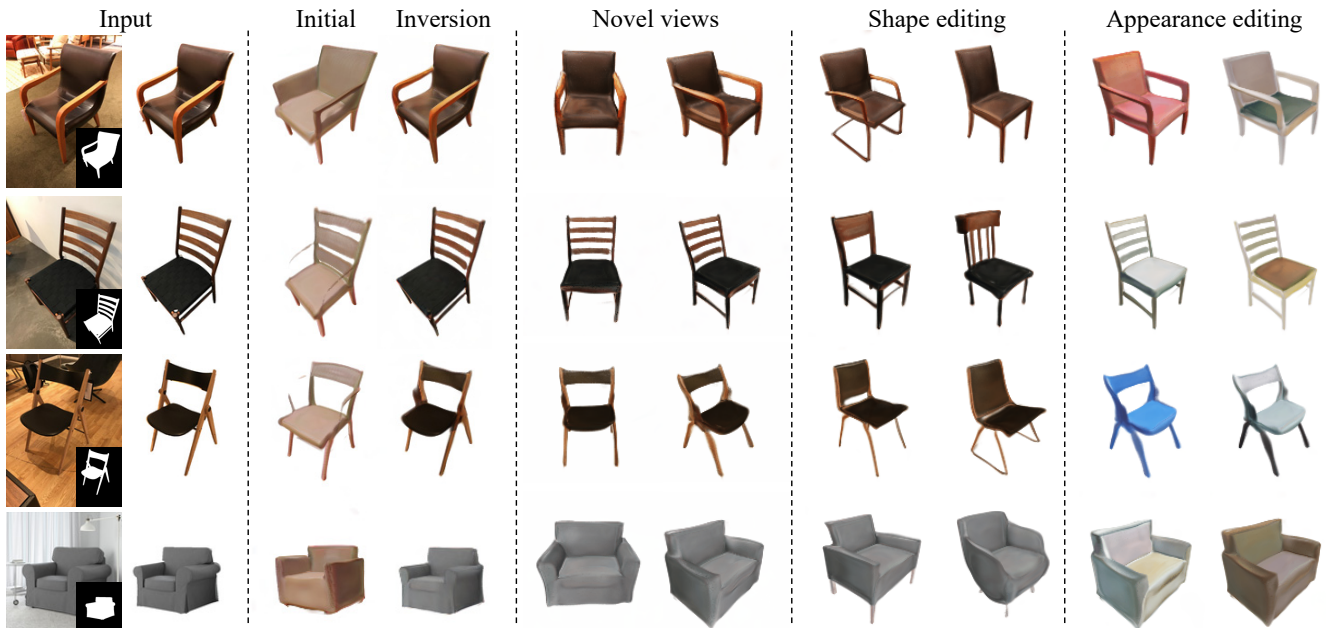


Figure 10: Real image inversion and disentangled editing results from our method.

Table 2: Ablation study on different network components and training losses on PhotoShape-chair 64×64 .

Setting	FID	KID ($\times 100$)
no deformation	24.6	1.28
no correction	18.4	0.59
w/o deformation smooth. loss	17.7	0.70
w/o rigidity loss	15.8	0.56
w/o normal consist. loss	18.6	0.82
w/o minimal correction loss	16.0	0.61
Full model	15.6	0.56

appearance code swapping experiments and compare with CodeNeRF [JA21] and EditNeRF [LZZ*21], both of which can separately control the shape and appearance attributes of an instance. Figure 7 shows the code swapping results of the generated CARLA objects. CodeNeRF and EditNeRF can hardly maintain a correct texture transfer result when swapping the appearance code between different instances (e.g., see the roof of the cars). By contrast, our method well keeps the texture consistency between corresponding semantic parts of different objects.

6.5. Real Image Inversion and Editing

Figure 10 shows some typical real image inversion and editing results by our method. Since our model is trained without background, we remove the irrelevant background content before inversion. As shown, our method can faithfully reconstruct the shape and appearance of the given real images and render plausible novel views of them, even though it is only trained with synthetic data. In addition, we can further edit the shape and appearance of the real images by replacing the inverted shape or appearance features with randomly generated ones.

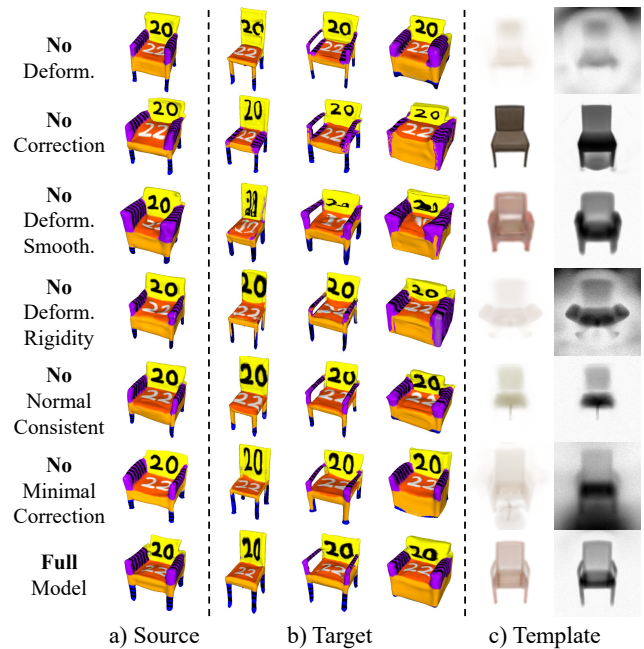


Figure 11: Visualization of the learned dense correspondence and template radiance field under different setups.

6.6. Ablation Study

We evaluate the efficacy of our deformable radiance field architecture and training regularization on PhotoShape chairs with 64×64 resolution. Table 2 shows the influence of different components on image generation quality. Our full model with deformation field, correction field, and all regularization losses yields the best result.

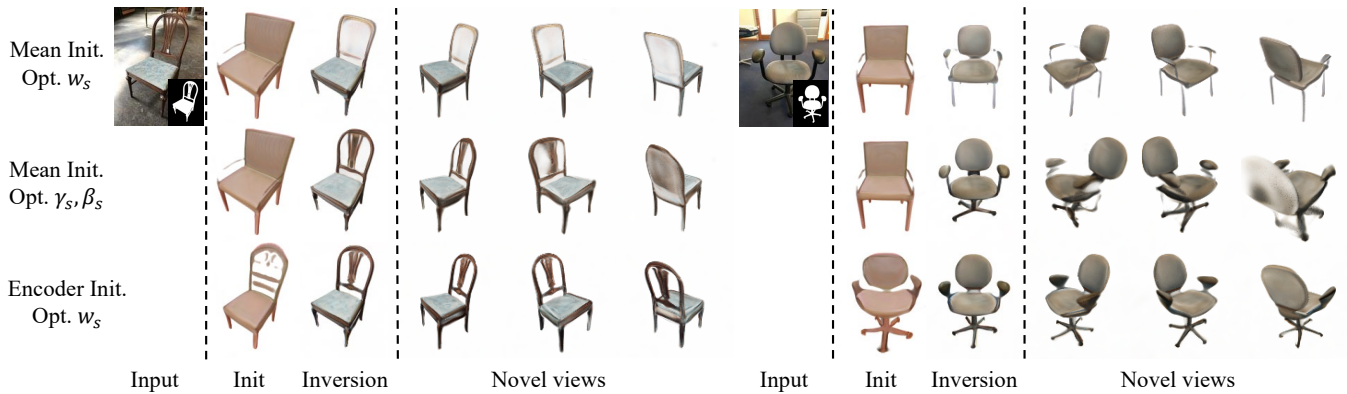


Figure 12: Comparison of different inversion initialization methods.

We further evaluate the influence of each component on the learned dense correspondence between different shapes in Fig. 11. To visualize dense correspondence, we extract isosurfaces from the density fields via MarchingCubes [LC87] and manually painted some textures on the extracted meshes, as shown in Fig. 11(a). With a textured mesh as source shape, we can transfer its appearance to other generated instances using the dense correspondence derived from the learned deformation field, and Fig. 11(b) shows the texture transfer results. We also visualize the learned template radiance fields in Fig. 11(c) where we render the corresponding color and depth images.

As shown in in Fig. 11, the correspondence is poor without the deformation or correction. The model without deformation can only derive correspondence by absolute spatial location, and without correction it is unable to separate parts with topology difference. The deformation smoothness and rigidity losses alleviate irregular deformation and local distortions. The normal consistency loss encourages the deformation field to learn a more precise correspondence, and the minimal correction loss prevents the density from being compensated mainly by the correction field and hence also leads to improved correspondence. Our full model ends up with a semantically informative template and yields the best quality in terms of dense correspondence.

To further verify the effectiveness of our proposed CNN-based inversion initialization, we compare it with the commonly-used mean latent code initialization strategy. Specifically, we first obtain a mean shape latent feature w_s^0 by averaging the features of 10K randomly generated instances. Then we test two settings: directly optimizing w_s with w_s^0 as initialization (i.e., same as ours defined in Eq. 15), or optimizing the latent features (γ_s, β_s) with the corresponding values derived from w_s^0 as initialization. As shown in Fig. 12, the former strategy leads to poor image reconstruct quality with inaccurate geometry and appearance. The latter one can well reconstruct the given image, but the geometry is erroneous and unnatural as can be seen from the novel view renderings. In contrast, our CNN-based initialization scheme can estimate a good starting point for the inversion process in the w_s space, which leads to our high-quality reconstruction and novel view rendering results.




Figure 13: Some failure cases of real image inversion from our trained model.

6.7. Running time

It takes our method 0.4 seconds to render a 128×128 image on a Nvidia Tesla V100 GPU. For image inversion, we use 2K optimization steps which takes about 15 minutes for a single instance. For our image editing experiments, we simply replace the shape or appearance features and hence they take the same amount of time as image rendering.

6.8. Failure Cases

We illustrate two failure cases of image inversion in Fig. 13. Our method may not be able to recover the correct geometry when the shapes in the input images are rare or the local structures are too complex. These cases are difficult to handle for both our trained generator and the inversion initialization CNN.

7. Conclusion

We have presented a novel generative deformable radiance field that can generate multiview images of an object category with topology variations, and achieve disentangled control over shape and appearance attributes of the synthesized images. The shape-appearance disentanglement is achieved in an unsupervised manner through our deliberate architecture and training loss design and no additional label is introduced compared to previous 3D-aware GAN training. We also demonstrated that our method can be applied to disentangled editing of real images by leveraging our proposed image inversion scheme. We believe that our method has strong potentials for realistic virtual content creations and manipulations in future applications.

References

- [AQW19] ABDAL R., QIN Y., WONKA P.: Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), pp. 4432–4441. 5
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: Scape: shape completion and animation of people. In *ACM SIGGRAPH 2005 Papers*. 2005, pp. 408–416. 3
- [AZMW21] ABDAL R., ZHU P., MITRA N. J., WONKA P.: Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics* 40, 3 (2021), 1–21. 3
- [BMT*21] BARRON J. T., MILDENHALL B., TANCIK M., HEDMAN P., MARTIN-BRUALLA R., SRINIVASAN P. P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *IEEE/CVF International Conference on Computer Vision* (2021). 2
- [BMV*22] BARRON J. T., MILDENHALL B., VERBIN D., SRINIVASAN P. P., HEDMAN P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022). 2
- [BSAG18] BIŃKOWSKI M., SUTHERLAND D. J., ARBEL M., GRETTON A.: Demystifying mmd gans. In *International Conference on Learning Representations* (2018). 8
- [BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3D faces. In *26th Annual Conference on Computer Graphics and Interactive Techniques* (1999), pp. 187–194. 2, 3
- [CDH*16] CHEN X., DUAN Y., HOUTHOOFT R., SCHULMAN J., SUTSKEVER I., ABBEEL P.: Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Advances in Neural Information Processing Systems* 29 (2016). 3
- [CFG*15] CHANG A. X., FUNKHOUSER T., GUIBAS L., HANRAHAN P., HUANG Q., LI Z., SAVARESE S., SAVVA M., SONG S., SU H., XIAO J., YI L., YU F.: *ShapeNet: An Information-Rich 3D Model Repository*. Tech. Rep. arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. 6
- [CLC*22] CHAN E. R., LIN C. Z., CHAN M. A., NAGANO K., PAN B., DE MELLO S., GALLO O., GUIBAS L., TREMBLAY J., KHAMIS S., ET AL.: Efficient geometry-aware 3d generative adversarial networks. In *IEEE/CVF International Conference on Computer Vision* (2022). 2
- [CMK*21] CHAN E. R., MONTEIRO M., KELLNHOFFER P., WU J., WETZSTEIN G.: pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 5799–5809. 1, 2, 3, 4, 5, 6, 8
- [DBS*21] DEVRIES T., BAUTISTA M. A., SRIVASTAVA N., TAYLOR G. W., SUSSKIND J. M.: Unconstrained scene generation with locally conditioned radiance fields. In *IEEE/CVF International Conference on Computer Vision* (2021). 2
- [DRC*17] DOSOVITSKIY A., ROS G., CODEVILLA F., LOPEZ A., KOLTUN V.: Carla: An open urban driving simulator. In *Conference on Robot Learning* (2017), pp. 1–16. 2, 6
- [DSTB16] DOSOVITSKIY A., SPRINGENBERG J. T., TATARCHENKO M., BROX T.: Learning to generate chairs, tables and cars with convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 4 (2016), 692–705. 2, 3
- [DYC*20] DENG Y., YANG J., CHEN D., WEN F., TONG X.: Disentangled and controllable face image generation via 3D imitative-contrastive learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 5154–5163. 3
- [DYT21] DENG Y., YANG J., TONG X.: Deformed implicit field: Modeling 3d shapes with learned dense correspondence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 10286–10296. 2, 3, 4, 5
- [DYXT22] DENG Y., YANG J., XIANG J., TONG X.: Gram: Generative radiance manifolds for 3d-aware image generation. In *IEEE/CVF International Conference on Computer Vision* (2022). 2, 4
- [GFK*18] GROUEIX T., FISHER M., KIM V. G., RUSSELL B. C., AUBRY M.: A papier-mâché approach to learning 3D surface generation. In *IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 216–224. 2, 3
- [HMP*16] HIGGINS I., MATTHEY L., PAL A., BURGESS C., GLOROT X., BOTVINICK M., MOHAMED S., LERCHNER A.: beta-vae: Learning basic visual concepts with a constrained variational framework. 3
- [HRU*17] HEUSEL M., RAMSAUER H., UNTERTHINER T., NESSLER B., HOCHREITER S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems* (2017), pp. 6626–6637. 8
- [IZZE17] ISOLA P., ZHU J.-Y., ZHOU T., EFROS A. A.: Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition* (2017), pp. 1125–1134. 2
- [JA21] JANG W., AGAPITO L.: Codenerf: Disentangled neural radiance fields for object categories. In *IEEE/CVF International Conference on Computer Vision* (2021), pp. 12949–12958. 2, 9
- [JHTG20] JIANG C., HUANG J., TAGLIASACCHI A., GUIBAS L.: Shapeflow: Learnable deformations among 3D shapes. In *Advances in Neural Information Processing Systems* (2020). 2
- [KLA19] KARRAS T., LAINE S., AILA T.: A style-based generator architecture for generative adversarial networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 4401–4410. 4
- [KLA*20] KARRAS T., LAINE S., AITTALA M., HELLSTEN J., LEHTINEN J., AILA T.: Analyzing and improving the image quality of stylegan. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 8110–8119. 5
- [KWKT15] KULKARNI T. D., WHITNEY W., KOHLI P., TENENBAUM J. B.: Deep convolutional inverse graphics network. In *Advances in Neural Information Processing Systems* (2015). 2
- [LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH* 21, 4 (1987), 163–169. 10
- [LMR*15] LOPER M., MAHMOOD N., ROMERO J., PONS-MOLL G., BLACK M. J.: SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics* 34, 6 (2015), 1–16. 2, 3
- [LSS*21] LOMBARDI S., SIMON T., SCHWARTZ G., ZOLLHOEFER M., SHEIKH Y., SARAGIH J.: Mixture of volumetric primitives for efficient neural rendering. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–13. 2
- [LZZ*21] LIU S., ZHANG X., ZHANG Z., ZHANG R., ZHU J.-Y., RUSSELL B.: Editing conditional radiance fields. In *IEEE/CVF International Conference on Computer Vision* (2021), pp. 5773–5783. 2, 9
- [MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989* (2022). 2
- [MGN18] MESCHEDER L., GEIGER A., NOWOZIN S.: Which training methods for gans do actually converge? In *International Conference on Machine Learning* (2018), pp. 3481–3490. 4
- [MON*19] MESCHEDER L., OECHSLE M., NIEMEYER M., NOWOZIN S., GEIGER A.: Occupancy networks: Learning 3d reconstruction in function space. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 4460–4470. 2
- [MST*20] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHY R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision* (2020), Springer, pp. 405–421. 1, 2, 4, 5
- [NG21] NIEMEYER M., GEIGER A.: Giraffe: Representing scenes as compositional generative neural feature fields. In *IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition* (2021), pp. 11453–11464. 2
- [NPLT*19] NGUYEN-PHUOC T., LI C., THEIS L., RICHARDT C., YANG Y.-L.: Hologan: Unsupervised learning of 3D representations from natural images. In *IEEE/CVF International Conference on Computer Vision* (2019), pp. 7588–7597. 2
- [PAM*18] PUMAROLA A., AGUDO A., MARTINEZ A. M., SANFELIU A., MORENO-NOGUER F.: Ganimation: Anatomically-aware facial animation from a single image. In *European conference on computer vision* (2018), pp. 818–833. 3
- [PFS*19] PARK J. J., FLORENCE P., STRAUB J., NEWCOMBE R., LOVEGROVE S.: DeepSDF: Learning continuous signed distance functions for shape representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 165–174. 2
- [PGM*19] PASZKE A., GROSS S., MASSA F., LERER A., BRADBURY J., CHANAN G., KILLEEN T., LIN Z., GIMELSHEIN N., ANTIGA L., ET AL.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019). 5
- [PKA*09] PAYSAN P., KNOTHE R., AMBERG B., ROMDHANI S., VETTER T.: A 3D face model for pose and illumination invariant face recognition. In *IEEE International Conference on Advanced Video and Signal Based Surveillance* (2009), pp. 296–301. 3
- [PLWZ19] PARK T., LIU M.-Y., WANG T.-C., ZHU J.-Y.: Semantic image synthesis with spatially-adaptive normalization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 2337–2346. 3
- [PNM*20] PENG S., NIEMEYER M., MESCHEDER L., POLLEFEYS M., GEIGER A.: Convolutional occupancy networks. In *European Conference on Computer Vision* (2020), Springer, pp. 523–540. 2
- [PRFS18] PARK K., REMATAS K., FARHADI A., SEITZ S. M.: Photo-shape: Photorealistic materials for large-scale shape collections. *ACM Transactions on Graphics* 37, 6 (Nov. 2018). 2, 6
- [PSB*21] PARK K., SINHA U., BARRON J. T., BOUAZIZ S., GOLDMAN D. B., SEITZ S. M., MARTIN-BRUALLA R.: Deformable neural radiance fields. In *IEEE/CVF International Conference on Computer Vision* (2021). 2
- [PZX*21] PENG S., ZHANG Y., XU Y., WANG Q., SHUAI Q., BAO H., ZHOU X.: Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 9054–9063. 2
- [SLNG20] SCHWARZ K., LIAO Y., NIEMEYER M., GEIGER A.: Graf: Generative radiance fields for 3d-aware image synthesis. In *Advances in Neural Information Processing Systems* (2020). 1, 2, 6
- [SLY*18] SHEN Y., LUO P., YAN J., WANG X., TANG X.: Faceid-gan: Learning a symmetry three-player gan for identity-preserving face synthesis. In *IEEE conference on computer vision and pattern recognition* (2018), pp. 821–830. 3
- [SMB*20] SITZMANN V., MARTEL J., BERGMAN A., LINDELL D., WETZSTEIN G.: Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems* 33 (2020). 2
- [STH*19] SITZMANN V., THIES J., HEIDE F., NIESSNER M., WETZSTEIN G., ZOLLHOFFER M.: Deepvoxels: Learning persistent 3d feature embeddings. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 2437–2446. 2
- [SZW19] SITZMANN V., ZOLLHÖFER M., WETZSTEIN G.: Scene representation networks: continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems* (2019), pp. 1121–1132. 2
- [TCY*22] TANCİK M., CASSER V., YAN X., PRADHAN S., MILDENHALL B., SRINIVASAN P. P., BARRON J. T., KRETZSCHMAR H.: Block-nerf: Scalable large scene neural view synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022). 2
- [TDB16] TATARCHENKO M., DOSOVITSKIY A., BROX T.: Multi-view 3d models from single images with a convolutional network. In *European Conference on Computer Vision* (2016), pp. 322–337. 2
- [TPF*22] TEWARI A., PAN X., FRIED O., AGRAWALA M., THEOBALT C., ET AL.: Disentangled3d: Learning a 3d generative model with disentangled geometry and appearance from monocular images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022). 2
- [TYL17] TRAN L., YIN X., LIU X.: Disentangled representation learning gan for pose-invariant face recognition. In *IEEE conference on computer vision and pattern recognition* (2017), pp. 1415–1424. 3
- [VBPP06] VLASIC D., BRAND M., PFISTER H., POPOVIC J.: Face transfer with multilinear models. In *ACM SIGGRAPH 2006 Courses*. 2006, pp. 24–es. 3
- [WCH*22] WANG C., CHAI M., HE M., CHEN D., LIAO J.: Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022). 1, 2
- [WLL*21] WANG P., LIU L., LIU Y., THEOBALT C., KOMURA T., WANG W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems* (2021). 2
- [WSK*15] WU Z., SONG S., KHOSLA A., YU F., ZHANG L., TANG X., XIAO J.: 3D ShapeNets: A deep representation for volumetric shapes. In *IEEE conference on Computer Vision and Pattern Recognition* (2015), pp. 1912–1920. 2
- [XPMBB21] XIE C., PARK K., MARTIN-BRUALLA R., BROWN M.: Fig-nerf: Figure-ground neural radiance fields for 3d object category modelling. In *2021 International Conference on 3D Vision (3DV)* (2021), IEEE, pp. 962–971. 2
- [XXH*21] XIANG F., XU Z., HASAN M., HOLD-GEOFFROY Y., SUNKAVALLI K., SU H.: Neutex: Neural texture mapping for volumetric neural rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 7119–7128. 2
- [YFKT*22] YU A., FRIDOVICH-KEIL S., TANCİK M., CHEN Q., RECHT B., KANAZAWA A.: Plenoxels: Radiance fields without neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022). 2
- [YFST18] YANG Y., FENG C., SHEN Y., TIAN D.: FoldingNet: Point cloud auto-encoder via deep grid deformation. In *IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 206–215. 3
- [YTB*21] YENAMANDRA T., TEWARI A., BERNARD F., SEIDEL H.-P., ELGHARIB M., CREMERS D., THEOBALT C.: i3dmm: Deep implicit 3d morphable model of human heads. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 12803–12813. 3
- [Zho21] ZHOU B.: Interpreting generative adversarial networks for interactive image generation. *arXiv preprint arXiv:2108.04896* (2021). 5
- [ZIE*18] ZHANG R., ISOLA P., EFROS A. A., SHECHTMAN E., WANG O.: The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), pp. 586–595. 5
- [ZKJB17] ZUFFI S., KANAZAWA A., JACOBS D. W., BLACK M. J.: 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2017), pp. 6365–6373. 3
- [ZSZZ20] ZHU J., SHEN Y., ZHAO D., ZHOU B.: In-domain gan inversion for real image editing. In *European conference on computer vision* (2020), Springer, pp. 592–608. 5
- [ZYDL21] ZHENG Z., YU T., DAI Q., LIU Y.: Deep implicit templates for 3D shape representation. In *IEEE Conference on Computer Vision and Pattern Recognition* (2021). 2, 3
- [ZYHC22] ZHENG M., YANG H., HUANG D., CHEN L.: Imface: A nonlinear 3d morphable face model with implicit neural representations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022). 3