

2D GANs Meet Unsupervised Single-view 3D Reconstruction

Feng Liu[✉], Xiaoming Liu[✉]

Michigan State University, Computer Science & Engineering
{liufeng6, liuxm}@msu.edu

Abstract. Recent research has shown that controllable image generation based on pre-trained GANs can benefit a wide range of computer vision tasks. However, less attention has been devoted to 3D vision tasks. In light of this, we propose a novel image-conditioned neural implicit field, which can leverage 2D supervisions from GAN-generated multi-view images and perform the single-view reconstruction of generic objects. Firstly, a novel offline StyleGAN-based generator is presented to generate plausible pseudo images with full control over the viewpoint. Then, we propose to utilize a neural implicit function, along with a differentiable renderer to learn 3D geometry from pseudo images with object masks and rough pose initializations. To further detect the unreliable supervisions, we introduce a novel uncertainty module to predict uncertainty maps, which remedy the negative effect of uncertain regions in pseudo images, leading to a better reconstruction performance. The effectiveness of our approach is demonstrated through superior single-view 3D reconstruction results of generic objects. Code is available at <http://cvlab.cse.msu.edu/project-gansvr.html>.

Keywords: 2D GANs, Multi-view Pseudo Images, Unsupervised, Single-view 3D Reconstruction, Generic objects, Uncertainty

1 Introduction

Realistic image synthesis is an important research area of computer vision. There has been remarkable progress in this field with the advent of 2D Generative Adversarial Networks (GANs) [9], such as StyleGAN [23] and its variations [21, 22, 24], which can generate high-fidelity images of diverse object categories with a wide variety of attributes (*e.g.*, pose, identity). Such superior capabilities of modeling the semantic image manifold enable the generated photorealistic images to be leveraged for many vision tasks, such as image editing [20, 55], domain translation [73], face recognition [30], and video generation [60]. However, it remains much less explored in 3D vision tasks, *e.g.*, 3D reconstruction [31, 32].

The 2D GAN manifolds appear to learn 3D geometrical properties implicitly, where recent GAN interpretation methods [14, 51] have shown that manipulating the latent code of the pre-trained GAN models can produce images of the same object under different viewpoints. Our work aims to answer the following




Fig. 1: Our approach leverages StyleGAN-generated multi-view pseudo images to learn a 3D model without 3D supervision, which can perform single-view 3D reconstruction for a variety of generic objects, *e.g.*, airplanes, birds, cars, horses, motorbikes, potted plants, *etc.* In addition, our framework produces uncertainty maps, indicating the unreliable local areas in the pseudo images.

question. *Using the GAN-generated multi-view images, can we learn a category-specific multi-view stereo system without 3D supervision that can reconstruct 3D shapes from a single image?*

Early attempts [44, 52, 63] are made to mine 3D geometric cues from the pre-trained 2D GAN models in an unsupervised manner. However, without modeling objects in the 3D space, these methods only recover 2.5D representations (depth or normals). Recently, StyleGANRender [71] integrates StyleGAN to generate multi-view images, which may be used to train an inverse graphics network for 3D reconstruction. However, the method focuses more on performing independent manipulation of 3D properties in GAN’s latent space by fine-tuning the GAN models. In addition, the unreliable texture existing in GAN-generated multi-view images is a common issue, and has not been investigated.

It remains a challenge to leverage GAN-generated multi-view supervision for single-view 3D reconstruction. First of all, the pre-trained 2D GAN models lack explicit and precise camera pose to control over generated images, which is a necessity for classic multi-view stereo. Second, the GAN-generated multi-view images often suffer from *local* distortion and low perceptual quality, which severely breaks the *consistency* of either object shape or texture across views, and thereby ruins the *cornerstone* of multi-view stereo.

To address these challenges, we propose a novel framework to leverage GAN-generated multi-view images (termed ‘*pseudo images*’) in learning generic object shape models, for the purpose of 3D reconstruction from a single image (Fig. 1). To first generate multi-view imagery by a pre-trained GAN, *e.g.*, StyleGAN, we carefully study the latent space of StyleGAN and devise a simple but effective technique, which generates plausible images with an azimuth range of 0–360°. Consequently, during training, given a realistic image generated by StyleGAN, we can produce a set of pseudo images of the same object under different view-

points. We then introduce a neural implicit network to simultaneously learn the unknown geometry, texture, and camera parameters for the objective of reconstructing the pseudo images, by incorporating a differentiable renderer.

A key component is that we introduce a learning framework that enables the neural implicit network to be conditioned on a single image. Specifically, we adopt an image encoder as a hypernetwork to predict the network parameters of the implicit function. This image conditioning allows the framework to be trained on multi-view images, where it learns object geometry priors within the category to perform single-view reconstruction. Moreover, to address the unreliable texture supervision issue in pseudo images, we devise an uncertainty prediction module, together with an uncertainty-aware photometric loss to estimate uncertainty maps, which can effectively filter out the unreliable supervision signals/inconsistencies within/across multi-view pseudo images, leading to a more precise reconstruction. Comprehensive experiments show the superiority of our method over existing methods in unsupervised single-view 3D reconstruction.

In summary, the contributions of this work include:

- ◊ We propose a novel image-conditioned neural implicit network, which can exploit 2D supervision from GAN-generated multi-view pseudo images and performs single-view 3D reconstruction of generic objects.
- ◊ We introduce a multi-view image generation mechanism based on the pre-trained StyleGAN models, which can produce plausible images with full control over viewpoints.
- ◊ We propose an uncertainty prediction module to ignore unreliable texture supervision in pseudo images, enabling a reliable self-supervised learning.
- ◊ Our method shows superior single-view 3D reconstruction for rigid and non-rigid generic objects in the wild.

2 Prior Work

Application of Pre-trained 2D GANs While research on GANs is rapidly growing, our review mainly focuses on the pre-trained unconditional 2D GAN models. The capability to produce high-quality images makes 2D GANs applicable to many vision tasks, *e.g.*, image restoration [61, 67], image editing (inpainting, super-resolution, semantic manipulation) [12, 45], segmentation [72], and DeepFake attack and defense [2, 6, 7, 48]. Further, the pre-trained 2D GAN models have been applied to data augmentation to reduce overfitting and bias in deep models [47, 54]. To expand to 3D vision applications, prior works [36, 56, 62, 63, 74] adopt GANs to learn 3D shapes from images but rely on either 3D supervision or a 3D generator, which suffers from heavy memory consumption or extra training difficulties. Recently, LiftedGAN [52] lifts a pre-trained StyleGAN and distill it into a 3D aware generator, producing depth maps as a by-product. Similarly, GAN2Shape [44] produces an unsupervised decomposition by using a GAN model as supervision. However, those methods require inefficient online image generation during training and infer 2.5D representations only. StyleGAN-Render [71] exploits StyleGAN as a multi-view generator to learn an *mesh-based*

Table 1: Comparison of *unsupervised* shape learning methods. [Keys: Cam. = camera poses per training sample, Requ. or Cons. = requirement or constraint, Real data= whether can train on real-world images, \otimes = camera poses for a set of reference images]

Method	Output Representation	Required Template	Required Cam.	Additional Requ. or Cons.	Real data
LiftedGAN [52]	2.5D, depth	\times	\times	GAN models, pre-trained	\checkmark
GAN2Shape [44]	2.5D, depth	\times	\times	GAN models, pre-trained	\checkmark
StyleGANRender [71]	3D, mesh	\times	\times	GAN models, fine-tuning	\checkmark
DVR [42]	3D, implicit	\times	\checkmark	multi-view	\times
DIST [34]	3D, implicit	\times	\checkmark	multi-view	\times
SDFDiff [17]	3D, implicit	\times	\checkmark	multi-view	\times
CSDM [58]	3D, mesh	\checkmark	\checkmark	2D semantic	\checkmark
CMR [18]	3D, mesh	\checkmark	\times	2D semantic	\checkmark
U-CMR [8]	3D, mesh	\checkmark	\times	viewpoint distribution	\checkmark
UMR [28]	3D, mesh	\times	\times	3D semantic	\checkmark
CSM [27]	3D, mesh	\checkmark	\times	-	\checkmark
A-CSM [26]	3D, mesh	\checkmark	\times	-	\checkmark
DRC [59]	3D, voxel	\times	\checkmark	multi-view	\checkmark
SRN [53]	3D, implicit	\times	\checkmark	multi-view	\times
NeRF [38]	3D, implicit	\times	\checkmark	multi-view	\checkmark
SDF-SRN [29]	3D, implicit	\times	\checkmark	-	\checkmark
ShSMesh [69]	3D, volumetric	\times	\times	-	\checkmark
Proposed	3D, implicit	\times	\otimes	GAN models, pre-trained	\checkmark

inverse graphics network to turn the StyleGAN into a controllable render. However, they require a *fine-tuning* step for the entire StyleGAN model, which is not desirable in this work. Moreover, they do not tackle the unreliability in pseudo multi-view images. In contrast, our method focuses on leveraging *pre-trained* 2D GANs for single-view 3D reconstruction. Despite both methods utilizing GAN-generated pseudo images for 3D modeling, we step forward in more plausible multi-view generation, robust shape and texture representation, and uncertainty-aware photometric supervision mechanism.

3D-aware Generative Models Understanding the latent representation of GANs has resulted in a body of works disentangling various factors of generated objects in a 3D-controllable manner, *e.g.*, viewpoint. These approaches can be classified into two groups. One adds additional modules or losses in training to explicitly disentangle 3D factors [37]. For example, HoloGAN [39] controls the object pose by rigid-body transformations via a 3D feature module. StyleFlow [1] learns non-linear paths in the latent space by normalizing flows conditioned on the attribute. Recently there has been a trend combining of the neural radiance fields (NeRF) [5, 11, 40, 41, 49] with GANs to devise 3D-aware generators. Another line of works, such as InterFaceGAN [50], SeFa [51], GANSpace [14], discover the latent semantic directions of a pre-trained GAN model that can manipulate object rotation unaware of its underlying 3D model. It is preferable to exploit the knowledge contained in a pre-trained GAN image manifold for the goal of recovering 3D object shapes without retraining the GAN models.

Shape Learning without 3D Supervision While 3D reconstruction, especially for faces [3, 33, 57], is a long-standing topic, we focus our review on shape learning from real-world images of generic objects without 3D supervision. Recent neural networks tackle this ill-posed problem via a differentiable renderer

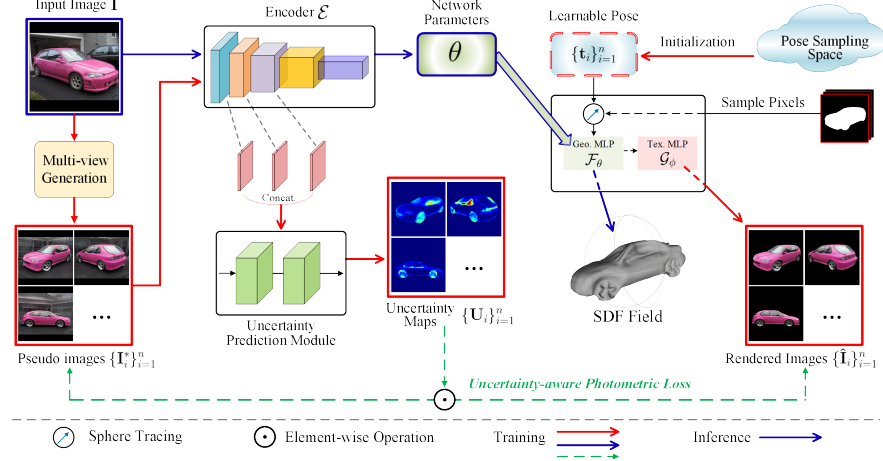


Fig. 2: **Overview.** The proposed framework is composed of two key modules: an offline StyleGAN-based multi-view generation and an image-conditioned neural implicit network. During training, the neural implicit network learns unknown geometry, texture, and camera poses for the objective of approximating the multi-view pseudo images. At inference time, the learned neural implicit function performs 3D reconstruction for the object from a single image.

along with a choice of 3D shape representation [17, 34, 35, 42, 53, 59]. However, in these works, multiple views of the same object with known cameras are required, which limits their learning from real-world images. Another branch of works show promising reconstruction from real-world images [8, 26–28, 64]. SDF-SRN [29] mines more supervision from 2D silhouette for superior reconstruction, yet still requires camera pose. ShSMesh [69] further discards the need for former constraints, but their reconstructions are of lower quality. NeRF [38] and its variations are scene- or object-specific models, which limit their applications for 3D reconstruction from unseen objects or scenes. In contrast, our models are category-specific, and can perform single-view 3D reconstruction for novel instances. Tab. 1 summarizes the differences between our method and prior work.

3 Proposed Method

We start with an overview of the proposed framework (Fig. 2) and then present the individual modules in detail. We first introduce an offline and effective multi-view generator based on the pre-trained StyleGAN models, which produce plausible multi-view images with full control over viewpoints. Then, we detail the proposed image-conditioned neural implicit field learning framework, including a neural implicit network, differentiable rendering procedure, and an uncertainty prediction module. These three modules work jointly for the objective of exploiting the pseudo images to learn generic object shape priors and perform 3D reconstruction from a single input image.

3.1 StyleGAN based Multi-view Generation

We briefly review the embedding space of the StyleGAN [23, 24]. Typically, a generator $G(\cdot)$ samples a latent code \mathbf{z} from a pre-defined distribution \mathcal{Z} such as the normal distribution, and produces an output image \mathbf{I} . The code \mathbf{z} is first mapped to an intermediate latent space \mathcal{W} via a Multi-layer Perceptron (MLP), and then \mathcal{W} is transformed to \mathcal{W}^+ space by 16 learned affine transformations. The generator $G(\cdot)$ projects \mathbf{W} to the final image: $G(\mathbf{W}) = \mathbf{I}$. Such latent codes have been shown to learn various disentangled semantics [14, 51]. For instance, StyleGAN-Render [71] finds that the latent codes $\mathbf{W}_v := (\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3, \mathbf{w}_4) \in \mathbb{R}^{4 \times 512}$ in the first 4 layers control camera viewpoints. That is, given a source and reference generated image pair $(\mathbf{I}^S, \mathbf{I}^R)$ with their latent codes $(\mathbf{W}^S, \mathbf{W}^R)$, we can generate an image of the source object with the reference viewpoint by swapping $(\mathbf{W}_v^S, \mathbf{W}_v^R)$ and keeping the rest dimensions of \mathbf{W}^S . We denote this multi-view generation strategy as **Baseline**. However, while \mathbf{W}_v indeed alters the object viewpoint, it still perceives the shape of the reference object, as shown in Fig. 3a.

It is difficult to develop a multi-view stereo system from noisy multi-view images with inconsistent shapes. To tackle this issue, inspired by SeFa [51], we propose a novel offline multi-view generation mechanism that generates plausible images with enhanced cross-view consistency. SeFa suggests that the weight parameters in early affine transformations contain essential knowledge of image variations. One can obtain interpretable directions in the latent space by computing eigenvectors of their weight matrices and selecting eigenvectors with the largest eigenvalues. With this observation, we propose to filter the viewpoint-irrelevant features in \mathbf{W}_v^R guided by the eigenvectors of the k largest eigenvalues, which are computed from the weight parameters in the first 4 transformations. It can be formulated as:

$$\arg \min_{\alpha} \|\hat{\mathbf{W}}_v - \mathbf{W}_v^R\|^2, \quad \hat{\mathbf{W}}_v = \alpha \mathbf{V} + \mathbf{W}_v^S, \quad (1)$$

where $\hat{\mathbf{W}}_v$ is the enhanced viewpoint latent code. α is a k -dim viewpoint coefficient. $\mathbf{V} \in \mathbb{R}^{k \times 512}$ denotes the eigenvectors of $\mathbf{A}^T \mathbf{A}$ associated with the k largest eigenvalues. $\mathbf{A} \in \mathbb{R}^{m \times 512}$ are the weights of transformations. $\alpha \mathbf{V}$ is duplicated into four rows in Eqn. 1 and α can be solved by gradient descent. Finally, we generate a new image by combining $\hat{\mathbf{W}}_v$ and remaining the dimensions of \mathbf{W}^S . **Training Pseudo Images** Given a pre-trained StyleGAN model, we first synthesize training images and filter out images that have more than one instance or an unrealistic instance, resulting in N training samples $\{\mathbf{I}^j\}_{j=1}^N$. Then, we manually select n reference view samples, which roughly cover the common object viewpoints ranging from $0 - 360^\circ$ in azimuth. Finally, for each training sample \mathbf{I}^j , we produce n multi-view images of the same object with fixed camera poses as the pseudo images $\{\mathbf{I}_i^{*j}\}_{i=1}^n \in \mathbb{R}^{W \times H \times 3}$. We further apply the work [46] to obtain corresponding instance segmentation $\{\mathbf{M}_i^j\}_{i=1}^n$ of pseudo images.

Pose Sampling Space We assume a pinhole camera model $\mathcal{C} = (\mathbf{t}, \mathbf{K})$, where $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ is the intrinsic parameter. We represent the camera pose/extrinsic parameters $\mathbf{t} = (\mathbf{c}_0, \mathbf{q})$ based on its 3D position $\mathbf{c}_0 \in \mathbb{R}^3$ and its rotation from a

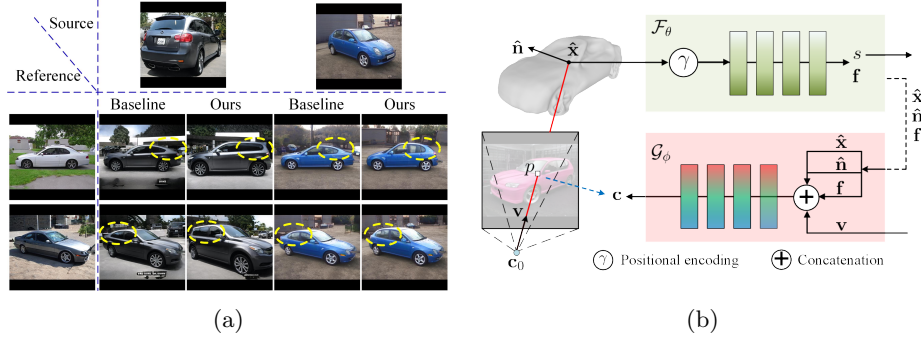


Fig. 3: (a) Multi-view generator comparisons of **Baseline** [71] and our approach (**Ours**). As can be observed, besides the viewpoint, the baseline perceives shape cues from the references (yellow circle, best view in zoom in). While our generated images show more consistency in object shape across views, which benefits shape learning. Please refer to **Supp** for more categories. (b) The architecture of our neural implicit fields and differentiable renderer.

canonical view. $\mathbf{q} \in \mathbb{R}^4$ is the quaternion vector representing the camera rotation. We assume the observed object is approximately inside the unit sphere. We further assume a known intrinsic and the principal point at the image center, as commonly assumed in stereo systems [19]. For the pose initialization, we manually annotate the camera poses $\{\tilde{\mathbf{t}}_i\}_{i=1}^n$ for the n reference images, which initializes the pseudo images' poses, $\mathbf{t}_i^j = \tilde{\mathbf{t}}_i, j \in [1, N]$. *Since we only need to annotate n reference samples per object category, it is far more practical than prior works that require camera pose label per training sample* (see Tab. 1).

3.2 Image-conditioned Neural Implicit Field

3D Geometry Representation As illustrated in Fig. 3b, the object geometry to be reconstructed is represented by the function [68]: $\mathcal{F}: \gamma(\mathbf{x}) \rightarrow (s, \mathbf{f})$ that maps a point $\mathbf{x} \in \mathbb{R}^3$ to its signed distance value s to the object surface and a local geometry feature $\mathbf{f} \in \mathbb{R}^{d_f}$. $\gamma(\cdot)$ denotes a positional encoding operator on \mathbf{x} with 6 exponentially increasing frequencies introduced in NeRF [38]. The surface \mathcal{S} is represented as the zero level set of MLP \mathcal{F} with learnable parameters θ :

$$\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^3 | \mathcal{F}_\theta^{(s)}(\gamma(\mathbf{x})) = 0\}. \quad (2)$$

Neural Renderer Given a pixel p of a masked input image, we march a ray $\mathbf{r} = \{\mathbf{c}_0 + t\mathbf{v} | t \geq 0\}$, where \mathbf{c}_0 is the camera position and \mathbf{v} the viewing direction. $\hat{\mathbf{x}}$ denotes the first intersection between the ray \mathbf{r} and the surface, which can be efficiently detected via the sphere tracing [15] and implemented in a differentiable manner. As in Fig. 3b, the rendered color of the pixel p is encoded as [68]

$$\mathcal{G}: (\hat{\mathbf{x}}, \hat{\mathbf{n}}, \mathbf{f}, \mathbf{v}) \rightarrow \mathbf{c}, \quad (3)$$

which is a function of the surface properties at $\hat{\mathbf{x}}_p$, including the surface point $\hat{\mathbf{x}}_p$, the surface normal $\hat{\mathbf{n}}_p \in \mathbb{R}^3$, the local geometry feature \mathbf{f}_p , and a viewing direction $\mathbf{v}_p \in \mathbb{R}^3$. Similarly, the function \mathcal{G} is implemented as an MLP with learnable parameters ϕ . The surface normal $\hat{\mathbf{n}}_p$ can be computed by the spatial derivative $\frac{\delta \mathcal{F}^{(s)}}{\delta \mathbf{x}_p}$ via back-propagation through the network \mathcal{F} . Incorporating the surface normal and view direction enables \mathcal{G} to represent the light reflected from a surface point \mathbf{x} across different viewpoints [38, 68]. Also, according to [68], introducing the local geometry feature vector \mathbf{f} allows the renderer to handle more complex appearances, which might appear in GAN-generated images. The neural implicit field can be learned through back-propagation without any 3D supervision by comparing the rendered images with the multi-view pseudo images. Minimizing this error encourages multi-view photo consistency because only when the point is on the actual surface, the MLPs (\mathcal{F}, \mathcal{G}) will predict accurate color with fewer multi-view variations.

Image Encoder We propose to utilize architecture to condition on a single image, such that the learned neural implicit field can generalize to a new object instance without re-training. Specifically, as shown in Fig. 2, we use an image encoder \mathcal{E} as a hyper network [13, 29] to predict θ (the parameters of \mathcal{F}), written as $\theta = \mathcal{E}_\Phi(\mathbf{I})$, where Φ is the neural network weights.

3.3 Model Learning

Given N sets of masked pseudo images $\{\mathbf{I}^j, \{\mathbf{I}_i^{*j}\}_{i=1}^n\}_{j=1}^N$, along with their initial camera poses $\{\{\mathbf{t}_i^j\}_{i=1}^n\}_{j=1}^N$, we optimize the encoder parameters Φ , texture MLP parameters ϕ and camera poses $\{\{\mathbf{t}_i^j\}_{i=1}^n\}_{j=1}^N$ by minimizing the loss:

$$\mathcal{L} = \sum_{j=1}^N (\mathcal{L}_{RGB} + \lambda_{mask} \mathcal{L}_{Mask} + \lambda_{eik} \mathcal{L}_{Eik}), \quad (4)$$

where \mathcal{L}_{RGB} is photometric loss, \mathcal{L}_{Mask} is silhouette loss, \mathcal{L}_{Eik} is Eikonal regularization, and λ_* are loss weights.

Photometric Loss Let $\mathbf{I}_p^*, \mathbf{M}_p \in \{0, 1\}$ be the RGB and silhouette values of pixel p in an image sample \mathbf{I}_i^* taken at the view direction \mathbf{v}_p associated with camera \mathcal{C}_i . $p \in P$ indexes all pixels in the input image set $\{\mathbf{I}_i^*\}_{i=1}^n$. The photometric loss is defined on mini-batches of pixels in P :

$$\mathcal{L}_{RGB} = \frac{1}{|P|} \sum_{p \in P^{in}} |\mathbf{I}_p^* - \mathcal{G}(\hat{\mathbf{x}}_p, \hat{\mathbf{n}}_p, \mathbf{f}_p, \mathbf{v}_p)|, \quad (5)$$

where $\mathbf{f}_p, \hat{\mathbf{x}}_p, \hat{\mathbf{n}}_p$ are defined in Eqn. 3. $P^{in} \subset P$ represents the subset of pixels P where intersection has been found and $\mathbf{M}_p = 1$. $|\cdot|$ denotes the L_1 loss.

Silhouette Loss We define the silhouette loss as

$$\mathcal{L}_{Mask} = \frac{1}{|P|} \sum_{p \in P^{out}} CE(\mathbf{M}_p, \hat{\mathbf{M}}_p), \quad (6)$$

where $\hat{\mathbf{M}}$ is the masked rendering. $P^{out} = P - P^{in}$ represents the indices in the mini-batch for which there is no ray-geometry intersection or $\mathbf{M}_p = 0$. $CE(\cdot, \cdot)$ denotes the cross-entropy loss. Conventionally, given the ray $\mathbf{r}_p = \{\mathbf{c}_0 + t\mathbf{v}_p | t \geq 0\}$ of pixel p , $\hat{\mathbf{M}}_p$ is defined as:

$$\hat{\mathbf{M}}_p = \begin{cases} 1 & \mathbf{r}_p \cap \mathcal{S} \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

To make this differentiable, we follow [68] and compute

$$\hat{\mathbf{M}}_p = \text{sigmoid}(-\beta \min_{t \geq 0} \mathcal{F}^{(s)}(\mathbf{c}_0 + t\mathbf{v}_p)). \quad (8)$$

When $\beta \rightarrow \infty$, $\mathcal{F}^{(s)} < 0$ means inside the surface and $\mathcal{F}^{(s)} > 0$ outside.

Eikonal Regularization A special property of signed distance functions is their differentiability with a gradient of unit norm, satisfying the Eikonal equation $\|\nabla \mathcal{F}\|_2 = 1$ [10, 43]. We thus encourage our implicit geometry representation to satisfy the Eikonal property:

$$\mathcal{L}_{Eik} = \sum_{\tilde{\mathbf{x}}} \left\| \|\nabla_{\tilde{\mathbf{x}}} \mathcal{F}^{(s)}(\tilde{\mathbf{x}})\|_2 - 1 \right\|_2^2, \quad (9)$$

where $\tilde{\mathbf{x}}$ is uniformly sampled at the 3D region of interest.

3.4 Uncertainty Prediction Module

The pseudo images are inherently aleatoric, *i.e.*, there might be areas with either notable artifacts in one image or with inconsistent shape/texture across images. To adaptively treat these problematic areas in learning, we propose to use Bayesian learning [25, 66] to model this aleatoric uncertainty. Specifically, we exploit the feature space of the image encoder \mathcal{E} and train a shallow decoder to estimate an uncertainty map \mathbf{U} , which has the same size as pseudo images. Formally, we model the observed color \mathbf{I}_p^* at pixel p with a likelihood function $\mathcal{P}(\mathbf{I}_p^*)$ that follows the Laplacian distribution with ray-dependent variance σ_p :

$$\mathcal{P}(\mathbf{I}_p^*) = \frac{1}{2\sigma_p} \exp\left(-\frac{|\mathbf{I}_p^* - \mathcal{G}(\hat{\mathbf{x}}_p, \hat{\mathbf{n}}_p, \mathbf{f}_p, \mathbf{v}_p)|}{\sigma_p}\right), \quad (10)$$

where σ_p denotes the uncertainty. Since L_1 distance is less sensitive to outliers, which is more suitable for optimizing the rendered appearance against the pseudo RGB values. Thus, we adopt Laplacian likelihood to model the inconsistent uncertainty distribution. To find the parameters best explaining the model, we maximize the likelihood function, *i.e.*, minimizing the negative log-likelihood:

$$-\log(\mathcal{P}(\mathbf{I}_p^*)) = \frac{|\mathbf{I}_p^* - \mathcal{G}(\hat{\mathbf{x}}_p, \hat{\mathbf{n}}_p, \mathbf{f}_p, \mathbf{v}_p)|}{\sigma_p} + \log\sigma_p + \log 2. \quad (11)$$

Therefore, we update the photometric loss as:

$$\mathcal{L}_{RGB} = \frac{1}{|P|} \sum_{p \in P^{in}} (e^{-\mathbf{U}_p} |\mathbf{I}_p^* - \mathcal{G}(\hat{\mathbf{x}}_p, \hat{\mathbf{n}}_p, \mathbf{f}_p, \mathbf{v}_p)| + \mathbf{U}_p). \quad (12)$$




Fig. 4: The uncertainty maps produced by our proposed method and humans.

We train this loss on mini-batches of pixels in P . Here P indexes all pixels in the input multi-view image set $\{\mathbf{I}_{i=1}^*\}_{i=1}^n$, which contribute to the same object’s depth, texture and uncertainty learning. During training, for the surface point \mathbf{x}_p , the first term $e^{-\mathbf{U}_p}$ can be seen as a weighted distance which assigns larger weights to less uncertain pixels. The second term \mathbf{U}_p is a penalty term.

\mathcal{L}_{RGB} encourages the neural implicit representation to bring together the texture information of all cross-view corresponding pixels in all images. Consequently, the pixel-wise uncertainty value is able to mine the multi-view inconsistencies among those pixels which contribute to the same 3D surface point. In practice, we train a 2-layer convolutional network to predict the log variance $\mathbf{U}_p := \log \sigma_p$ (please refer to **Supp** for more details).

3.5 Implementation Details

The encoder \mathcal{E} is implemented as a ResNet-18 [16] followed by fully-connected layers. Both \mathcal{F} and \mathcal{G} have 4 fully-connected layers. For the main experiment, we set $N=2,000$, $n=40$, $k=5$, $W=H=256$, $d_f=256$, $\beta=50$, $\lambda_{mask}=0.01$, $\lambda_{eik}=0.1$. We implement our model in Pytorch and use the Adam optimizer with a learning rate of $1e-4$ for both network and camera pose parameters.

4 Experimental Results

We evaluate our approach on six category-specific StyleGAN models, including rigid objects such as airplanes, cars, motorbikes, as well as non-rigid objects such as birds, horses, and potted plants. We use the official car and horse models from StyleGAN2 repo ¹, trained on the LSUN dataset [70]. For other 4 categories, we train the StyleGAN models with StyleGAN2-ADA-Pytorch library ² on LSUN airplanes, birds, motorbikes, and potted plants with 200K images respectively.

4.1 Human Study vs Our Method on the Uncertainty Prediction

Although GAN interpretation methods have shown that manipulating the latent code of StyleGAN produces multi-view images of the same object [14, 51], no studies have quantitatively evaluated the unreliable/inconsistent object shape

¹ <https://github.com/NVlabs/stylegan2>

² <https://github.com/NVlabs/stylegan2-ada-pytorch>

or texture in/across the multi-view pseudo images. Thanks to our multi-view-stereo-like neural implicit network, our uncertainty map can serve as a means to detect the unreliable/inconsistent areas in the GAN-generated multi-view images. On the other hand, volunteers were asked to label the potentially problematic areas in the pseudo images. As shown in Fig. 4, given an image, the multi-view generator is able to generate pseudo images with varying viewpoints. We believe humans are able to reason unreliable/inconsistent regions in/across the pseudo images. Specifically, this is accomplished using a random set of 100 images from the PASCAL3D+ car category. For each image, we generate a pseudo image with a different viewpoint (Sec. 3.1). Then, the volunteers manually labels polygon-based regions of interest (uncertainty region) on the pseudo images, using the Matlab Image Labeler app. As can be observed in Fig. 4, human labels mainly focus on the global object shape inconsistency across views, which might not have the granularity to evaluate the pixel-level inconsistency. Nevertheless, we quantify the detection ability of our uncertainty maps by using the human labels as the ground-truth. We achieve 34.6% Intersection over Union (IoU), which shows our capability in detecting the inconsistent object shape in GAN-generated multi-view pseudo images. To our knowledge, this is the first method that tries to investigate the unreliable supervision across the GAN-generated pseudo multi-view images in 3D object modeling.

4.2 Quantitative 3D Reconstruction Evaluation

We quantitatively evaluate on the PASCAL3D+ dataset [65], a 3D reconstruction benchmark of real-world images with (approximate) CAD model annotations. Similar to prior work [29, 59], we use annotations of airplane and car categories on the test set for evaluation.

Evaluation Metrics. We adopt standard 3D reconstruction metrics: IoU and Chamfer- L_1 Distance (CD). Following [59], we compute 3D IoU between ground truth and prediction with the resolution of 32^3 . Following [29], we uniformly sample 3D points from the ground truth and prediction to compute CD.

Baselines We compare against SoTA unsupervised single-view 3D reconstruction baselines: CSDM [58], DRC [59], CMR [18], U-CMR [8] and SDF-SRN [29]. As detailed in Tabs. 1 and 2a, some baselines require additional unsupervised constraints, *e.g.*, DRC (implicit) and SDF-SRN (implicit) both require ground-truth camera pose for each training sample. The mesh-based methods such as CSDM, CMR, and U-CMR need expert object-specific templates as additional constraints. Here, we do not compare with ShSMesh [69] as it neither quantitatively evaluates on PASCAL3D+, nor trains on real-world car/airplane images. Also, we do not quantitatively compare with StyleGANRender [71] as the code or trained model is not publicly available, and StyleGANRender does not report full 3D shape reconstruction errors as our baselines did.

Results We present the comparisons of our approach in Tab. 2a and visualize sample predictions in Fig. 5. It can be observed that the **Proposed** model is significantly better than baselines in both CD (10.1% relative over SDF-SRN) and IoU (7.8% relative over DRC). It is worth noting that our models trained with

Table 2: (a) Quantitative 3D reconstruction results on PASCAL3D+. During training, CSDM, DRC and SDF-SRN require ground-truth camera pose per training sample, CMR uses 2D keypoints and object-specific templates as additional constraints, and U-CMR only relies on object-specific templates. [Keys: Requ. or Cons.=requirement or constraint in training, T=category-specific templates, C=poses per training sample, C*=poses for reference images, K=2D keypoints, S=semantic information, GANs=pre-trained GAN models]. All CD values are scaled by 10 following [29]. (b) Ablation of uncertainty prediction and the number of pseudo images, using the PASCAL3D+ car category.

(a)					
Category	Requ. or Cons.	Airplane		Car	
		CD (\downarrow)	IoU (\uparrow)	CD (\downarrow)	IoU (\uparrow)
CSDM [58]	T, C	—	0.400	—	0.600
DRC [59]	C	—	0.420	—	0.670
CMR [18]	T, K	0.625	—	0.474	0.640
U-CMR [8]	T	—	—	—	0.646
UMR [28]	S	—	—	—	0.620
SDF-SRN [29]	C	0.303	0.405	0.233	0.653
SDF-SRN* [29]	C	0.297	0.412	0.230	0.661
Proposed	GANs, C*	0.286	0.473	0.195	0.702

(b)					
	Proposed w/o Uncertainty ($n=40$)	Proposed ($n=$)			
		10	20	40	50
CD (\downarrow)	0.208	0.336	0.243	0.195	0.191
IoU (\uparrow)	0.681	0.612	0.643	0.702	0.714

Proposed setting only require ground-truth pose annotations for the reference images. It is more practical for real-world scenarios than DRC and SDF-SRN, which require ground-truth camera pose per training sample. Further, we re-train SDF-SRN with our GAN-generated training data and report the results in Tab. 2a (SDF-SRN*). It can be observed, despite the minor improvement over the original SDF-SRN due to our pseudo images, the new model still performs worse than ours. Fig. 5 shows visual comparisons to SDF-SRN [29] and StyleGANRender [71] results. As can be observed, our approach suffers slightly from shape ambiguity, *e.g.*, windows of cars tend to be concave. Nonetheless, our predictions more closely resemble the ground truth.

4.3 Ablation Study

All ablations use the models trained on the car category.

Effect on Uncertainty Prediction We evaluate single-view 3D reconstruction on a model trained without uncertainty prediction, *i.e.* using Eqn. 5 instead of Eqn. 12. As shown in Tab. 2b, our uncertainty prediction module can remedy the negative impact of uncertain texture in GAN-generated multi-view images, leading to improved 3D reconstruction (CD: 0.208 \rightarrow 0.195).

Effect on n The key assumption, as well as motivation of our work, is that GAN-generated multi-view images can be leveraged to learn a multi-view stereo system. To validate the impact of the amount of pseudo images, we train models with different numbers of image viewpoints, $n = 10, 20, 40, 50$. Tab. 2b shows that the model trained with $n = 40$ images significantly outperforms the ones with $n = 10, 20$, and saturates when $n = 40 \rightarrow 50$. Considering the tradeoff between reconstruction accuracy and training cost, we use $n = 40$ for all experiments.

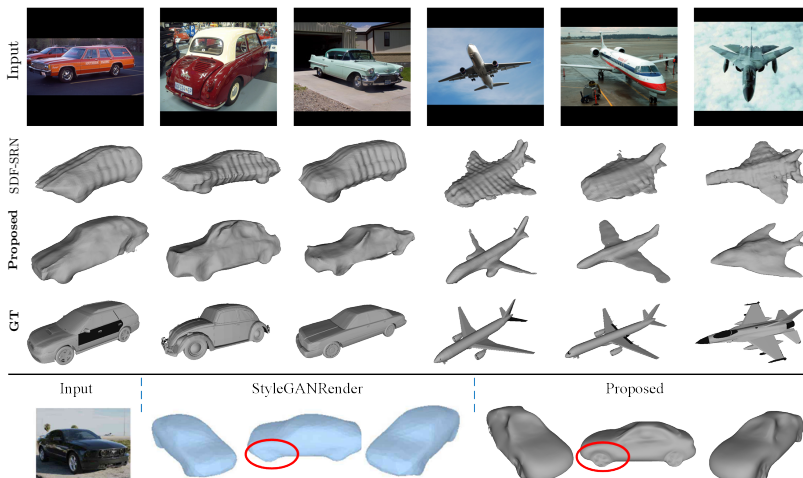


Fig. 5: Qualitative comparisons with SDF-SRN [29] (SOTA baseline) and StyleGANRender [71] on PASCAL3D+ car or airplane categories. Our approach recovers significantly more accurate 3D shapes and topologies from the images.

Effect on Multi-view Generation Following the same setting, we re-train a model with pseudo images produced by *Baseline* method (Sec. 3.1). Quantitatively, such a model only obtains the IoU of 0.679, much worse than ours (0.702). The comparisons show the superior quality of our multi-view generation method.

4.4 Qualitative Evaluation

Comparison with U-CMR [8] and DRC [59] We show qualitative comparisons with U-CMR and DRC on the PASCAL3D+ cars, motorbikes or airplanes, in Fig. 6. Our approach achieves more faithful reconstructions than U-CMR. Note that U-CMR requires expert templates while our approach does not.

Results on More Categories While our quantitative evaluation is on the PASCAL3D+ airplane and car, we provide more qualitative results for the birds, horses, motorbikes and potted plants in Fig. 7. As can be seen, our approach can effectively capture the thin structure present in 3D shapes from single-view images, *e.g.* horses’ legs and motorbikes’ hand clutch. All testing images are from the LSUN dataset, which never appear in the training set of our GAN models.

5 Conclusions

To leverage pre-trained GAN models for 3D vision tasks, we propose an image-conditioned neural implicit network that can learn the shape priors from GAN-generated multi-view images and perform single-view 3D reconstruction. Moreover, we naturally introduce a novel uncertainty prediction module, which can

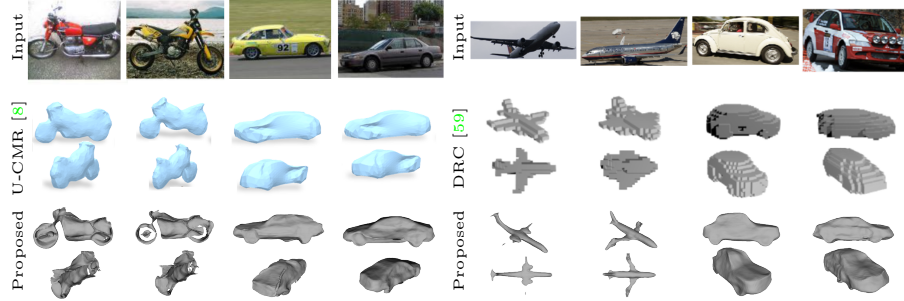


Fig. 6: Additional qualitative comparisons with U-CMR [8] and DRC [59] on cars and motorcycles or airplanes.

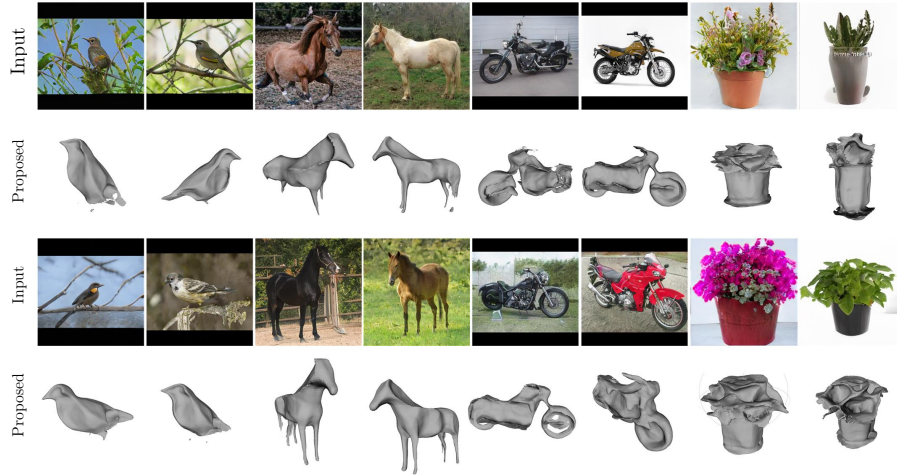


Fig. 7: Qualitative results of birds, horses, motorbikes, and potted plants.

avoid invalid supervisions for better single-view 3D reconstruction. Experimentally, our approach significantly outperforms the SoTA unsupervised single-view 3D reconstruction methods, while requiring less supervision during training. We believe this work opens up a path for improving the ability to semantically control GAN generation and facilitates 2D GAN priors for 3D vision tasks.

Limitations For some categories (*e.g.*, chair), StyleGAN is unable to converge to satisfying results, partially due to chairs’ large topology variations. We believe the rapid development of GANs will extend to these challenging categories, and thus our method can leverage them for 3D reconstruction of more categories. Also, similar to most prior works, our model is category-specific. One future direction is to develop a single model for multiple categories based on universal GAN models, *e.g.*, BigGAN [4], improving generalization to unseen categories.

References

1. Abdal, R., Zhu, P., Mitra, N.J., Wonka, P.: StyleFlow: Attribute-conditioned exploration of styleGAN-generated images using conditional continuous normalizing flows. *TOG* (2021) 4
2. Asnani, V., Yin, X., Hassner, T., Liu, X.: Reverse engineering of generative models: Inferring model hyperparameters from generated images. *arXiv:2106.07873* (2021) 3
3. Bai, Z., Cui, Z., Rahim, J.A., Liu, X., Tan, P.: Deep facial non-rigid multi-view stereo. In: *CVPR* (2020) 4
4. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: *ICLR* (2019) 14
5. Chan, E.R., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-GAN: Periodic implicit generative adversarial networks for 3D-aware image synthesis. In: *CVPR* (2021) 4
6. Dang, H., Liu, F., Stehouwer, J., Liu, X., Jain, A.K.: On the detection of digital face manipulation. In: *CVPR* (2020) 3
7. Dolhansky, B., Bitton, J., Pflaum, B., Lu, J., Howes, R., Wang, M., Ferrer, C.C.: The deepfake detection challenge (DFDC) dataset. *arXiv preprint arXiv:2006.07397* (2020) 3
8. Goel, S., Kanazawa, A., Malik, J.: Shape and viewpoint without keypoints. In: *ECCV* (2020) 4, 5, 11, 12, 13, 14
9. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *NeurIPS* (2014) 1
10. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099* (2020) 9
11. Gu, J., Liu, L., Wang, P., Theobalt, C.: StyleNeRF: A style-based 3d-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985* (2021) 4
12. Gu, J., Shen, Y., Zhou, B.: Image processing using multi-code GAN prior. In: *CVPR* (2020) 3
13. Ha, D., Dai, A., Le, Q.V.: Hypernetworks. *arXiv preprint arXiv:1609.09106* (2016) 8
14. Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: GANSpace: Discovering interpretable GAN controls. In: *NeurIPS* (2020) 1, 4, 6, 10
15. Hart, J.C.: Sphere tracing: A geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer* (1996) 7
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR* (2016) 10
17. Jiang, Y., Ji, D., Han, Z., Zwicker, M.: SDFDiff: Differentiable rendering of signed distance fields for 3D shape optimization. In: *CVPR* (2020) 4, 5
18. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning category-specific mesh reconstruction from image collections. In: *ECCV* (2018) 4, 11, 12
19. Kar, A., Häne, C., Malik, J.: Learning a multi-view stereo machine. In: *NeurIPS* (2017) 7
20. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: *ICLR* (2018) 1
21. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. In: *NeurIPS* (2020) 1

22. Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. In: NeurIPS (2021) [1](#)
23. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019) [1](#), [6](#)
24. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: CVPR (2020) [1](#), [6](#)
25. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: NeurIPS (2017) [9](#)
26. Kulkarni, N., Gupta, A., Fouhey, D.F., Tulsiani, S.: Articulation-aware canonical surface mapping. In: CVPR (2020) [4](#), [5](#)
27. Kulkarni, N., Gupta, A., Tulsiani, S.: Canonical surface mapping via geometric cycle consistency. In: ICCV (2019) [4](#), [5](#)
28. Li, X., Liu, S., Kim, K., De Mello, S., Jampani, V., Yang, M.H., Kautz, J.: Self-supervised single-view 3D reconstruction via semantic consistency. In: ECCV (2020) [4](#), [5](#), [12](#)
29. Lin, C.H., Wang, C., Lucey, S.: SDF-SRN: Learning signed distance 3D object reconstruction from static images. In: NeurIPS (2020) [4](#), [5](#), [8](#), [11](#), [12](#), [13](#)
30. Liu, F., Kim, M., Jain, A., Liu, X.: Controllable and guided face synthesis for unconstrained face recognition. In: ECCV (2022) [1](#)
31. Liu, F., Liu, X.: Voxel-based 3D detection and reconstruction of multiple objects from a single image. NeurIPS (2021) [1](#)
32. Liu, F., Tran, L., Liu, X.: Fully understanding generic objects: Modeling, segmentation, and reconstruction. In: CVPR (2021) [1](#)
33. Liu, F., Zhao, Q., Liu, X., Zeng, D.: Joint face alignment and 3d face reconstruction with application to face recognition. TPAMI (2018) [4](#)
34. Liu, S., Zhang, Y., Peng, S., Shi, B., Pollefeys, M., Cui, Z.: DIST: Rendering deep implicit signed distance function with differentiable sphere tracing. In: CVPR (2020) [4](#), [5](#)
35. Liu, S., Chen, W., Li, T., Li, H.: Soft Rasterizer: Differentiable rendering for unsupervised single-view mesh reconstruction. In: ICCV (2019) [5](#)
36. Lunz, S., Li, Y., Fitzgibbon, A., Kushman, N.: Inverse graphics GAN: Learning to generate 3D shapes from unstructured 2D data. In: NeurIPS (2020) [3](#)
37. Medin, S.C., Egger, B., Cherian, A., Wang, Y., Tenenbaum, J.B., Liu, X., Marks, T.K.: MOST-GAN: 3D morphable stylegan for disentangled face image manipulation. In: AAAI (2022) [4](#)
38. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020) [4](#), [5](#), [7](#), [8](#)
39. Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., Yang, Y.L.: HoloGAN: Unsupervised learning of 3D representations from natural images. In: ICCV (2019) [4](#)
40. Niemeyer, M., Geiger, A.: Campari: Camera-aware decomposed generative neural radiance fields. arXiv preprint arXiv:2103.17269 (2021) [4](#)
41. Niemeyer, M., Geiger, A.: GIRAFFE: Representing scenes as compositional generative neural feature fields. In: CVPR (2021) [4](#)
42. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In: CVPR (2020) [4](#), [5](#)
43. Osher, S., Fedkiw, R., Piechor, K.: Level set methods and dynamic implicit surfaces. Appl. Mech. Rev. (2004) [9](#)

44. Pan, X., Dai, B., Liu, Z., Loy, C.C., Luo, P.: Do 2D GANs know 3D shape? unsupervised 3D shape reconstruction from 2D image GANs. In: ICLR (2021) [2](#), [3](#), [4](#)
45. Pan, X., Zhan, X., Dai, B., Lin, D., Loy, C.C., Luo, P.: Exploiting deep generative prior for versatile image restoration and manipulation. TPAMI (2021) [3](#)
46. Qin, X., Zhang, Z., Huang, C., Dehghan, M., Zaiane, O.R., Jagersand, M.: U2-Net: Going deeper with nested u-structure for salient object detection. Pattern Recognition (2020) [6](#)
47. Rojtberg, P., Pöllabauer, T., Kuijper, A.: Style-transfer GANs for bridging the domain gap in synthetic pose estimator training. In: AIVR (2020) [3](#)
48. Rossler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Nießner, M.: Face-forensics++: Learning to detect manipulated facial images. In: ICCV (2019) [3](#)
49. Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: GRAF: Generative radiance fields for 3D-aware image synthesis. In: NeurIPS (2020) [4](#)
50. Shen, Y., Yang, C., Tang, X., Zhou, B.: InterfaceGAN: Interpreting the disentangled face representation learned by GANs. TPAMI (2020) [4](#)
51. Shen, Y., Zhou, B.: Closed-form factorization of latent semantics in GANs. In: CVPR (2021) [1](#), [4](#), [6](#), [10](#)
52. Shi, Y., Aggarwal, D., Jain, A.K.: Lifting 2D StyleGAN for 3D-Aware face generation. In: CVPR (2021) [2](#), [3](#), [4](#)
53. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3D-structure-aware neural scene representations. In: NeurIPS (2019) [4](#), [5](#)
54. Su, K., Zhou, E., Sun, X., Wang, C., Yu, D., Luo, X.: Pre-trained StyleGAN based data augmentation for small sample brain CT motion artifacts detection. In: ADMA (2020) [3](#)
55. Suzuki, R., Koyama, M., Miyato, T., Yonetsuji, T., Zhu, H.: Spatially controllable image synthesis with internal representation collaging. arXiv preprint arXiv:1811.10153 (2018) [1](#)
56. Szabó, A., Meishvili, G., Favaro, P.: Unsupervised generative 3D shape learning from natural images. arXiv preprint arXiv:1910.00287 (2019) [3](#)
57. Tran, L., Liu, X.: On learning 3d face morphable model from in-the-wild images. TPAMI (2019) [4](#)
58. Tulsiani, S., Kar, A., Carreira, J., Malik, J.: Learning category-specific deformable 3D models for object reconstruction. TPAMI (2016) [4](#), [11](#), [12](#)
59. Tulsiani, S., Zhou, T., Efros, A.A., Malik, J.: Multi-view supervision for single-view reconstruction via differentiable ray consistency. In: CVPR (2017) [4](#), [5](#), [11](#), [12](#), [13](#), [14](#)
60. Tulyakov, S., Liu, M.Y., Yang, X., Kautz, J.: MoCoGAN: Decomposing motion and content for video generation. In: CVPR (2018) [1](#)
61. Wang, X., Li, Y., Zhang, H., Shan, Y.: Towards real-world blind face restoration with generative facial prior. In: CVPR (2021) [3](#)
62. Wu, J., Zhang, C., Xue, T., Freeman, W.T., Tenenbaum, J.B.: Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In: NeurIPS (2016) [3](#)
63. Wu, S., Rupprecht, C., Vedaldi, A.: Unsupervised learning of probably symmetric deformable 3D objects from images in the wild. In: CVPR (2020) [2](#), [3](#)
64. Wu, Y., Sun, Z., Song, Y., Sun, Y., Zhong, Y., Shi, J.: Shape-pose ambiguity in learning 3D reconstruction from images. In: AAAI (2021) [5](#)
65. Xiang, Y., Mottaghi, R., Savarese, S.: Beyond pascal: A benchmark for 3D object detection in the wild. In: WACV (2014) [11](#)

66. Xu, H., Zhou, Z., Wang, Y., Kang, W., Sun, B., Li, H., Qiao, Y.: Digging into uncertainty in self-supervised multi-view stereo. In: ICCV (2021) [9](#)
67. Yang, T., Ren, P., Xie, X., Zhang, L.: GAN prior embedded network for blind face restoration in the wild. In: CVPR (2021) [3](#)
68. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Basri, R., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. In: NeurIPS (2020) [7](#), [8](#), [9](#)
69. Ye, Y., Tulsiani, S., Gupta, A.: Shelf-supervised mesh prediction in the wild. In: CVPR (2021) [4](#), [5](#), [11](#)
70. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. arXiv preprint arXiv:1506.03365 (2015) [10](#)
71. Zhang, Y., Chen, W., Ling, H., Gao, J., Zhang, Y., Torralba, A., Fidler, S.: Image GANs meet differentiable rendering for inverse graphics and interpretable 3D neural rendering. In: ICLR (2021) [2](#), [3](#), [4](#), [6](#), [7](#), [11](#), [12](#), [13](#)
72. Zhang, Y., Ling, H., Gao, J., Yin, K., Lafleche, J.F., Barriuso, A., Torralba, A., Fidler, S.: DatasetGAN: Efficient labeled data factory with minimal human effort. In: CVPR (2021) [3](#)
73. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017) [1](#)
74. Zhu, J.Y., Zhang, Z., Zhang, C., Wu, J., Torralba, A., Tenenbaum, J.B., Freeman, W.T.: Visual object networks: Image generation with disentangled 3D representation. In: NeurIPS (2018) [3](#)