

Texture Fields: Learning Texture Representations in Function Space

Michael Oechsle^{1,2} Lars Mescheder¹ Michael Niemeyer¹ Thilo Strauss^{2†} Andreas Geiger¹
¹Autonomous Vision Group, MPI for Intelligent Systems and University of Tübingen
²ETAS GmbH, Bosch Group, Stuttgart

¹{firstname.lastname}@tue.mpg.de

[†]{firstname.lastname}@etas.com

Abstract

In recent years, substantial progress has been achieved in learning-based reconstruction of 3D objects. At the same time, generative models were proposed that can generate highly realistic images. However, despite this success in these closely related tasks, texture reconstruction of 3D objects has received little attention from the research community and state-of-the-art methods are either limited to comparably low resolution or constrained experimental setups. A major reason for these limitations is that common representations of texture are inefficient or hard to interface for modern deep learning techniques. In this paper, we propose *Texture Fields*, a novel texture representation which is based on regressing a continuous 3D function parameterized with a neural network. Our approach circumvents limiting factors like shape discretization and parameterization, as the proposed texture representation is independent of the shape representation of the 3D object. We show that *Texture Fields* are able to represent high frequency texture and naturally blend with modern deep learning techniques. Experimentally, we find that *Texture Fields* compare favorably to state-of-the-art methods for conditional texture reconstruction of 3D objects and enable learning of probabilistic generative models for texturing unseen 3D models. We believe that *Texture Fields* will become an important building block for the next generation of generative 3D models.

1. Introduction

3D reconstruction is one of the grand goals of computer vision. Recently, the vision community has witnessed impressive progress in reconstruction tasks like single image 3D reconstruction [6, 7, 9, 13, 23, 25, 43, 45] and generating 3D objects [4, 32, 46] using learning-based techniques which resolve ambiguities by incorporating prior knowledge. However, previous work on learning-based 3D reconstruction has mainly focused on the problem of reconstructing geometry. In contrast, texture reconstruction of 3D objects has received less attention from the community.

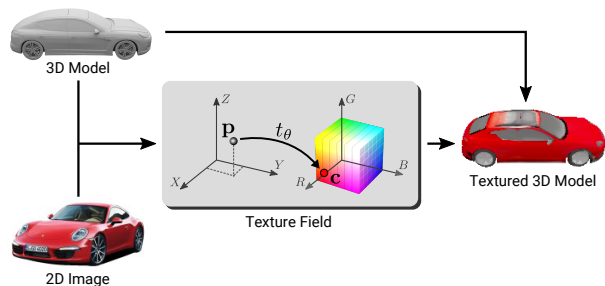


Figure 1: **Texture Fields** take a 3D shape and (optionally) a single 2D image of an object as input and learn a continuous function t_θ which maps any 3D point \mathbf{p} to a color value \mathbf{c} , enabling the prediction of accurately textured 3D models.

Previous approaches for texture reconstruction from single images differ in how texture information is represented. Several recent works [39, 42] have proposed volumetric voxel representations for colored 3D reconstruction. Unfortunately, however, due to the computational cost of voxel representations, which grows cubically with the resolution, the voxelizations are limited to rather low resolutions (usually 32^3 or 64^3) and hence cannot represent high frequency details. An alternative representation of texture consists of a 2D texture atlas and a parameterized mesh using a UV-mapping that maps a point on the shape manifold to a pixel in the texture atlas. However, current approaches based on mesh representations usually assume known topology and a predefined template mesh which limits these approaches to specific object categories such as birds [18] or faces [8, 36]. Predicting texture in the general case for arbitrary shapes without template mesh remains an unsolved problem.

Contributions: The main limiting factor of the aforementioned methods is their texture representation which either does not allow for high resolution reconstruction or strongly relies on task-specific shape parameterizations, limiting generality of these methods. An ideal texture representation, in contrast, should be independent of the shape representations (voxels, point cloud, mesh, etc.) and be able to represent high frequency detail. Towards this goal, we

propose Texture Fields as a novel representation of texture. Our key idea is to learn a continuous function for representing texture information in 3D space, see Fig. 1. By parameterizing this function through a deep neural network we are able to integrate this representation into a deep learning pipeline for 3D texture reconstruction that can be trained end-to-end.¹

Our experiments on various 3D object categories demonstrate that Texture Fields are able to predict high frequency texture information from a single image. We combine our approach with a state-of-the-art 3D reconstruction method [25], yielding a method which jointly reconstructs both the 3D shape and the texture of an object. Besides conditional experiments, we also extend our novel texture representation to the generative setting and show preliminary results for texture synthesis given a 3D shape model and a latent texture code. We conduct experiments on texture transfer between different objects as well as interpolations in the latent texture space which demonstrate that our generative model learns a useful representation of texture.

2. Related Work

We now briefly discuss the most related works on single image 3D reconstruction, generative image modeling using 3D knowledge and continuous representation learning.

2.1. Single Image Reconstruction

In recent years, 3D reconstruction of shapes from single images has made great progress [6, 7, 9, 13, 23, 25, 43, 45].

Voxels: Tulsiani et al. [42] proposed a voxel-based texture representation for learning colored 3D reconstruction based on ray consistency and multi-view supervision. More recently, Sun et al. [39] combined 2D-to-3D appearance flow estimation and voxel color regression for learning to reconstruct colored voxelizations in a supervised fashion. Unfortunately, however, voxel-based representations are severely limited in resolution due to computational and memory constraints. In contrast, our continuous approach does not require discretization and thus results in more details compared to voxel-based representations as we show in our experimental evaluation in Section 4.

Point Clouds: In recent years, novel view synthesis [27, 31, 40, 48] has been used for reconstructing colored point clouds from single images. By combining novel view synthesis and depth estimation, the method proposed in [40] reprojects the predicted image into a colored point cloud. In [27], point clouds are reconstructed from a set of novel views using multi-view stereo algorithms. Unfortunately, point-cloud based representations are sparse. While dense mesh representation can be extracted from point clouds, the

reconstructed shapes typically do not match the quality of state-of-the-art 3D reconstruction approaches and inferring the missing texture information requires additional post-processing steps. In contrary, our approach allows for inferring appearance for any location in 3D space and can be used in combination with arbitrary shape representations.

Meshes: Mesh-based approaches rely on category-specific template models and rigid texture parameterizations [8, 18, 36]. In contrast, our reconstruction approach can represent texture for arbitrary meshes without requiring a UV texture map of a category-specific template model. Note that determining a proper UV-mapping for arbitrary mesh models is a non-trivial problem which is typically solved with various heuristics for atlas generation. The advantage of our approach is that we circumvent mesh parameterization by disentangling the texture from the shape representation.

2.2. Generative Models

Recent work has shown that image generation methods [3, 11, 19, 20, 24, 30, 47] can be improved by exploiting 3D knowledge about the generated shapes [1, 2, 49]. Alhajja et al. [2] propose a model which learns to translate intrinsic properties (depth, normals, material) into RGB images. Zhu et al. [49] learn to predict 3D geometry as well as texture information in 2D images by disentangling shape, texture and pose in an unsupervised setting. In contrast to those image-based approaches, we directly predict the texture for the entire object in 3D.

2.3. Continuous Representations

Recently, parameterized continuous functions gained popularity for 3D shape reconstruction. Several works [6, 17, 25, 28] proposed to formulate 3D reconstruction as learning a continuous occupancy function or a signed-distance field in 3D space, parameterized via a neural network. Furthermore, in [10], a continuous function for color values was used for generating 2D images. In this work, we extend the concept of learning a representation in function space to reconstruct and generate texture information of 3D objects.

3. Method

This section describes the proposed Texture Field representation and demonstrates how it can be applied to conditional and unconditional texture synthesis tasks.

3.1. Texture Fields

Recent approaches to 3D reconstruction [6, 25, 28] represent 3D shapes as continuous functions of occupancy probability or surface distance. In contrast to point-, voxel- or mesh-based representations, these approaches do not rely on a fixed discretization and thus form an ideal basis for

¹Also see [35] for concurrent work that proposes a similar idea.

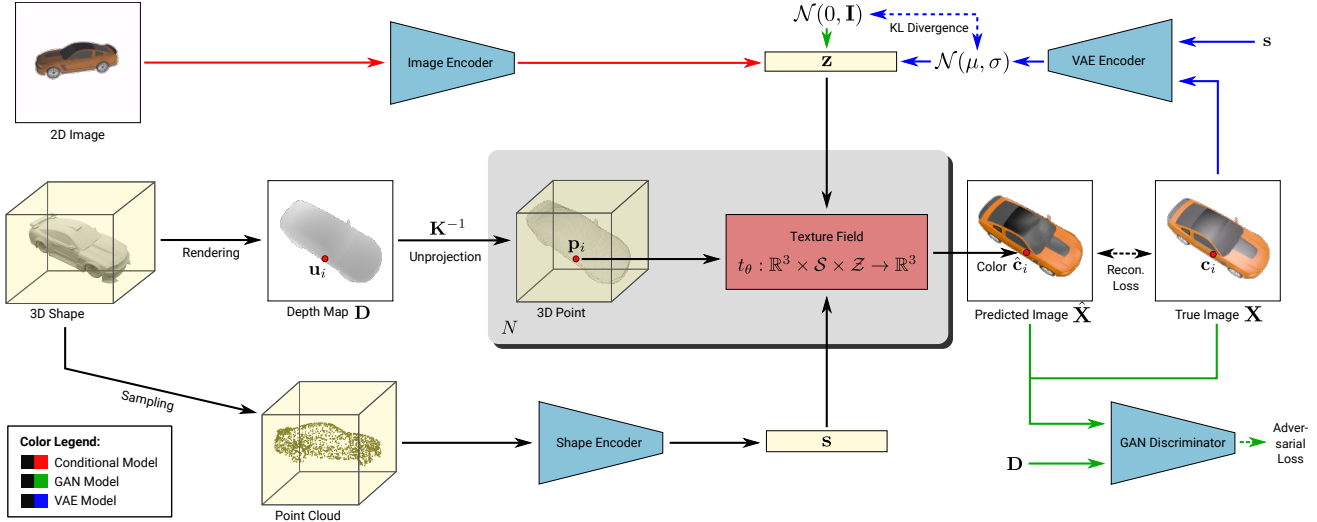


Figure 2: **Model Overview.** Colored arrows show *alternative* pathways representing the **Conditional**, **GAN** and **VAE** model. The blue and red boxes denote trainable components of our model which are parameterized through neural networks, respectively. The 3D shape of the object is encoded into a fixed-length vector representation s . Additionally, we render a depth map D from a randomly chosen viewpoint and obtain the corresponding 3D points \mathbf{p}_i by unprojecting all N image pixels \mathbf{u}_i into 3D. The reconstruction loss minimizes the difference between the pixel colors predicted by the Texture Field $\hat{c}_i = t_\theta(\mathbf{p}_i, s, \mathbf{z})$ and the ground truth pixel colors c_i . For training a **Conditional** model, the latent variable \mathbf{z} encodes information about the input image. In the unconditional case (i.e., for the **GAN** and **VAE**), \mathbf{z} is sampled from a Gaussian distribution.

representing appearance information. We explore this idea by embedding surface texture as a continuous function in 3D space. In combination with state-of-the-art 3D reconstruction techniques this allows us to reconstruct a textured 3D model from a single 2D image.

Let $t(\cdot)$ denote a function which maps a 3D point $\mathbf{p} \in \mathbb{R}^3$ to a color value $\mathbf{c} \in \mathbb{R}^3$ hence representing texture information by means of a 3D vector field:

$$t : \mathbb{R}^3 \rightarrow \mathbb{R}^3 \quad (1)$$

By parameterizing the function t as a neural network t_θ with learnable parameters θ , we reduce the problem to a simple regression task. However, the problem of texture generation remains ill-posed without any further constraints. Thus, in order to inform t_θ about the input shape of the object, we condition t_θ on a shape embedding $s \in \mathcal{S}$. This helps the Texture Field to guide its predictions towards the object surface and allows for exploiting contextual geometric information such as surface discontinuities which are often aligned with image edges. As input shape, we explore 3D CAD models as well as image-based 3D reconstructions using neural networks [25] in this paper.

Unfortunately, even given the input 3D shape, there still exists a variety of plausible texture explanations. Consider cars, for instance, where the 3D geometry alone does neither determine the color, nor the exact shape of the windows or headlights. However, we may further constrain the task by providing information about the object appearance.

More specifically, we condition t_θ on an additional 2D image taken from an arbitrary viewpoint. Note that an image provides only partial appearance information as it only depicts the object from a single perspective. Furthermore, we encode the image into a viewpoint-invariant global feature representation $\mathbf{z} \in \mathcal{Z}$. Thus, our approach does not assume the camera extrinsics wrt. the object to be known and therefore is able to texture untextured 3D shapes using images “in the wild”. Importantly, the input image need not depict an object of the exact same shape as the 3D model. This would be a strong restriction in practice as often only an approximate shape can be retrieved from a single image.

In summary, we define a *Texture Field* as a mapping from 3D point \mathbf{p} , shape embedding s and condition \mathbf{z} to color \mathbf{c} :

$$t_\theta : \mathbb{R}^3 \times \mathcal{S} \times \mathcal{Z} \rightarrow \mathbb{R}^3 \quad (2)$$

In the following, we will consider both the conditional case as well as the unconditional case. For the unconditional case, we exploit probabilistic generative networks [11, 22], capturing ambiguity in a *random* latent code \mathbf{z} .

Model: An overview over our *Texture Field* model is shown in Fig. 2. Colored arrows indicate alternative pathways representing the **Conditional**, **GAN** and **VAE** model. Blue and red boxes denote the trainable components of our model which are parameterized by neural networks whose parameters are trained jointly. We now provide details on each of the components of our model.

Shape Encoder: To inform the Texture Field t_θ about the 3D shape of the object, we uniformly sample 3D points from the input shape (typically provided in form of a triangular mesh) and pass the resulting point cloud to a PointNet encoder [29], adopting the network architecture of [25]. This results in a fixed-dimensional shape embedding \mathbf{s} .

Image Encoder: For the conditional case (red arrows in Fig. 2), we additionally extract appearance information from an input image. More specifically, we encode the input image into a fixed-dimensional latent code \mathbf{z} using a standard pre-trained residual network [14] with 18 layers.

Texture Field: Given the input shape \mathbf{s} and the condition \mathbf{z} , our proposed Texture Field model is able to predict a color value \mathbf{c}_i for any 3D point \mathbf{p}_i . Thus, it would be possible to directly color every point on the 3D mesh. Unfortunately, mesh-based appearance representations require additional UV-mappings for representing texture. We therefore train our texture model in 2D image space which provides a regular and hence efficient representation.²

Towards this goal, we render depth maps \mathbf{D} and corresponding color images \mathbf{X} from arbitrary viewpoints using OpenGL. The color at pixel \mathbf{u}_i and depth d_i is predicted as

$$\hat{\mathbf{c}}_i = t_\theta(d_i \mathbf{R} \mathbf{K}^{-1} \mathbf{u}_i + \mathbf{t}, \mathbf{s}, \mathbf{z}) \quad (3)$$

where i denotes the index for pixels with finite depth values $i \in \{1, \dots, N\}$ and N refers to the number of foreground pixels in the image (i.e., pixels where the object is visible). Here, the camera intrinsics and extrinsics are denoted by $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ and $(\mathbf{R} \in \mathbb{R}^{3 \times 3}, \mathbf{t} \in \mathbb{R}^3)$, respectively, and pixel \mathbf{u}_i is represented in homogeneous coordinates. For training our model, we compare the predicted color $\hat{\mathbf{c}}_i$ to the color of the corresponding pixel \mathbf{c}_i in the rendered image \mathbf{X} .

3.2. Training

This sections describes how we train our conditional and unconditional models. See Fig. 2 for a visual illustration.

Conditional Setting: In the conditional case, we input an embedding \mathbf{z} of the input image to our network.³ We train $t_\theta(\mathbf{p}, \mathbf{s}, \mathbf{z})$ in a supervised setting by minimizing the ℓ_1 -loss between the predicted image $\hat{\mathbf{X}}$ and the rendered image \mathbf{X} :

$$\mathcal{L}_{cond} = \frac{1}{B} \sum_{b=1}^B \sum_{i=1}^{N_b} \|t_\theta(\mathbf{p}_{bi}, \mathbf{s}_b, \mathbf{z}_b) - \mathbf{c}_{bi}\|_1 \quad (4)$$

Here, B represents the mini-batch size. Each element of the mini batch represents an image with N_b foreground pixels

²Note that this does not restrict our model in any way. At test time, our model can be evaluated for arbitrary 3D points.

³It is important to note that the input image and the images \mathbf{X} used for supervising the model during training need *not* be the same as we learn a viewpoint-invariant representation \mathbf{z} .

(i.e., pixels where the object is visible). Note that the shape encoding \mathbf{s}_b and the input encoding \mathbf{z}_b implicitly depend on the parameters of the shape and image encoder networks, respectively. We train the parameters of all three networks (shape encoder, image encoder, Texture Field) jointly.

Unconditional Generative Model: In the unconditional setting we are only given the 3D shape as input but no additional information about the appearance of the object. As mentioned above, this is a highly ill-posed task with many *valid* explanations. We therefore utilize probabilistic generative models which capture the ambiguity in the output using a latent code \mathbf{z} sampled from a Gaussian distribution (see green and blue models in Fig. 2). In particular, we adapt two recent deep latent variable models to our setting: a generative adversarial networks (GAN) [11] and a variational auto-encoder (VAE) [22]. Both models have been applied to a variety of image-based tasks in the past [11, 12, 19, 22, 30, 33] and are thus ideally suited in the context of our image-based loss formulation.

Let us first consider adversarial training. The problem of learning a generative model for texture information given a 3D shape can be tackled using a *conditional GAN* [26] where the generator is conditioned on the 3D shape. The generator is represented as a Texture Field $t_\theta : \mathbb{R}^3 \times \mathcal{S} \times \mathcal{Z} \rightarrow \mathbb{R}^3$ which maps the latent code \mathbf{z} for every given 3D location \mathbf{p}_i conditioned on the shape embedding \mathbf{s} to an RGB image:

$$\hat{\mathbf{X}} = G_\theta(\mathbf{z}_b | \mathbf{D}_b, \mathbf{s}_b) = \{t_\theta(\mathbf{p}_{bi}, \mathbf{s}_b, \mathbf{z}_b) | i \in \{1, \dots, N_b\}\} \quad (5)$$

As above, b denotes one element of the mini-batch and \mathbf{s}_b depends on the parameters of the shape encoder. We use a standard image-based discriminator $D_\phi(\mathbf{X}_b | \mathbf{D}_b)$ conditioned on the input depth image \mathbf{D} by concatenating it with the input image. For training the model, we use a non-saturating GAN loss with R_1 -regularization [24].

An alternative method for learning a latent variable model is given by a *conditional VAE* (cVAE) [37]. Our cVAE model comprises an encoder network that maps color image \mathbf{X} to mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\sigma}$ of an isotropic normal distributed random variable \mathbf{z} which follows the distribution $q_\phi(\mathbf{z} | \mathbf{X}, \mathbf{s})$. Our Texture Field model is used as decoder by predicting the color value $\hat{\mathbf{c}}_i$ of each pixel \mathbf{u}_i for the corresponding 3D locations \mathbf{p}_i conditioned on the shape embedding \mathbf{s} . Following [22, 33], we minimize the variational lower bound

$$\mathcal{L}_{VAE} = \frac{1}{B} \sum_{b=1}^B \left[\beta \text{KL}(q_\phi(\mathbf{z} | \mathbf{X}_b, \mathbf{s}_b) \| p_0(\mathbf{z}_b)) + \sum_{i=1}^{N_b} \|t_\theta(\mathbf{p}_{bi}, \mathbf{s}_b, \mathbf{z}_b) - \mathbf{c}_{bi}\|_1 \right] \quad (6)$$

where KL refers to the Kullback-Leibler divergence,

$p_0(\mathbf{z}) = \mathcal{N}(\mathbf{z}|0, \mathbf{I})$ denotes the standard normal distribution, \mathbf{z}_b is a sample from the posterior distribution $q_\phi(\mathbf{z}|\mathbf{X}_b, \mathbf{s}_b)$ and β is a trade-off parameter between the KL-divergence and the reconstruction loss [16]. Again, b denotes one element of the mini-batch and \mathbf{s}_b depends on the parameters of the shape distribution. In practice, we set $\beta = 1$. During training, we optimize \mathcal{L}_{VAE} using the reparameterization trick [22, 33]. At test time, we obtain new texture samples for the 3D object under consideration by sampling \mathbf{z} from a standard normal distribution.

3.3. Implementation details

In all of our experiments, we implement the texture field $t_\theta(\cdot, \mathbf{s}, \mathbf{z})$ using the fully connected ResNet [14] architecture from [25], see supplementary for details. For the image encoder, we use a ResNet-18 architecture [14], pre-trained on ImageNet. For the shape encoder, we adopt the PointNet [29] architecture from [25]. The GAN discriminator and the VAE encoder are based on the discriminator from [24]. We train both our supervised model and our VAE model end-to-end using Adam [21] with learning rate 10^{-4} . Our GAN is trained with alternating gradient descent using the RMSProp Optimizer [41] with the same learning rate.

4. Experimental Evaluation

We evaluate our approach in three different experiments. In the first part, we investigate the **representation power** of Texture Fields by analyzing how well the method can represent high frequency textures when trained on a single 3D object. In the second part, we apply our method to the challenging task of **single view texture reconstruction** in which we predict the full texture of 3D objects given only the 3D shape and a single view of this object. Moreover, we combine our method with a state-of-the-art shape reconstruction method [25] for full textured 3D reconstruction which also includes reconstructing the shape of the 3D object. In the last experiment, we explore if our representation can also be used in a **generative setting** where the goal is to produce a varied distribution of possible textures given only the shape of the 3D object without further input.

Baselines: Only few prior works have considered texture reconstruction in the general case without predefined template mesh or information about the camera view. Therefore, we construct a first simple baseline by mapping the input image onto the mesh by projecting all vertices into the input view to determine their color. While this simple baseline receives additional information about the camera of the input view, we believe it still serves as a good sanity check to see if our method is actually learning something useful and if it is able to correctly fill-in occluded regions. As a second baseline we consider a novel-view-synthesis (NVS) approach which uses the same image encoder as our ap-

proach, but applies a UNet architecture⁴ [34] to transform a depth rendering of the object into an RGB image. While this approach can also generate novel views of the object under consideration, it requires additional (lossy) post-processing to generate a complete texture map of the object. In particular, there is no guarantee for this baseline that the newly generated views are consistent under viewpoint changes. Lastly, we consider Im2Avatar [39] as a baseline for full textured 3D reconstruction from single images. To the best of our knowledge, this method currently achieves the highest resolution (64^3) among all voxel-based 3D reconstruction methods which predict colors and produces state-of-the-art results. We use the official implementation⁵ of Im2Avatar.

Dataset: Unless specified otherwise, we use the categories ‘cars’, ‘chairs’, ‘airplanes’, ‘tables’, ‘benches’, ‘cabinets’, ‘sofas’, ‘lamps’ and ‘vessels’ from the ShapeNet dataset [5] as these categories contain rich texture information while also providing a large variety of shapes. Note that all models in the dataset are axis-aligned. For our conditional experiments, we use the renderings provided by Choy et al. [7] as input. For training our models, we render 10 additional images and depth maps per object from random views in the upper hemisphere.

Metrics: For evaluation, we consider three different metrics in image space for random views of the objects. Firstly, to evaluate how well our method and the baselines capture the appearance distribution for a given object category, we use the Fréchet Inception Distance (FID) [15]. This is a common metric between distributions of images and widely used in the GAN literature [3, 15, 19, 20]. Moreover, to better estimate the distance between a predicted view and the ground truth on a per-instance basis, we measure the structure similarity image metric (SSIM) [44]. Since we found that SSIM mostly captures local properties of images, we additionally introduce a more global perceptual measure, which we call *Feature- ℓ_1 -metric*. Similarly to FID, the Feature- ℓ_1 -metric is computed by embedding both the image produced by the method under consideration and the ground truth view into feature space using an Inception network. The Feature- ℓ_1 -metric then computes the mean absolute distance between the feature activations of the predicted and the ground truth image.

4.1. Representation Power

The goal of this experiment is to explore the representation power of our Texture Field model. This “overfitting” experiment helps to disentangle the quality of the image encoder from the Texture Field representation itself and provides an upper bound on the reconstruction quality that we can expect when applying Texture Fields to more difficult

⁴See supplementary material for details.

⁵<https://github.com/syb7573330/im2avatar>



Ground Truth Texture Field Voxelization (128^3)

Figure 3: **Representation Power.** Comparison of a Texture Field fitted to a 3D model of a cat/human to a voxel-based representation and the corresponding ground truth model.

tasks. We train our method separately on 3D meshes of a cat and a human⁶, wrt. 512^2 px renderings from 500 views.

Qualitative results are shown in Fig. 3, comparing our model to a voxelization at a fixed resolution of 128^3 voxels. We observe that Texture Fields can represent high-frequency information while the voxel representation is inherently restricted to a Manhattan world of limited granularity. As we show in the supplementary material, Texture Fields can also represent textures of multiple object at once. These experiments validate that Texture Fields are a promising texture representation.

4.2. Single Image Texture Reconstruction

We now turn our attention to the challenging task of single-image texture reconstruction. Towards this goal, we conduct experiments for texture reconstruction given a 3D model together with a 2D image of the object from a random camera view. We train our method using the conditional setup described in Section 3.2. At test time we apply the trained model in three different settings. In the first setting, we use ground truth 3D shapes as input and apply our model to synthetic renderings of the object. In the second setting, we combine our method with a state-of-the-art shape reconstruction approach [25] to obtain a full 3D reconstruction pipeline of both shape and textures. Finally, in

⁶3D models from free3d.com and turbosquid.com.

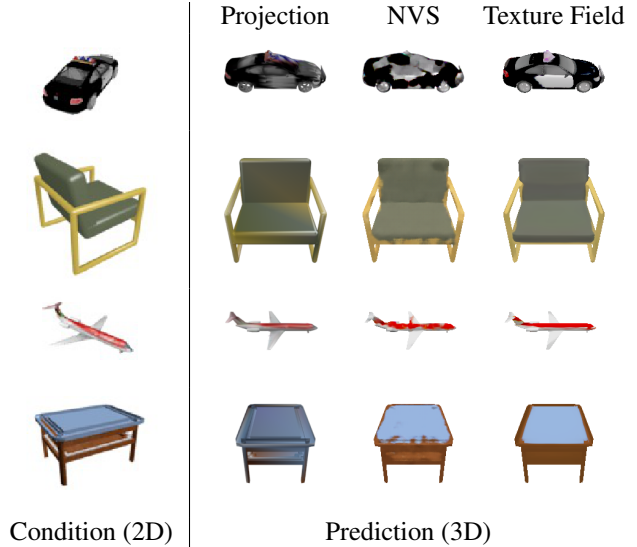


Figure 4: **Texture Reconstruction.** A qualitative comparison between our method (Texture Field) and the baselines is shown. In this experiment, we use GT shapes as input.

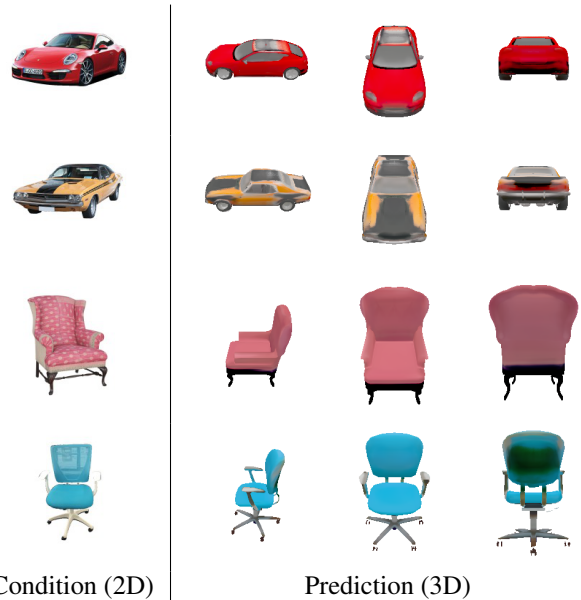


Figure 5: **Texture Reconstruction from Real Images.** In this experiment, our model transfers texture from previously unseen *real* images to unseen 3D CAD models.

the third setting, we investigate if our model also transfers to real data. Here, we use images of real cars together with similar 3D shapes from the ShapeNet dataset.

GT Shapes: When using ground truth shapes, our model successfully learns to extract texture from just a single image of the object as illustrated in Fig. 4. In particular, our model can complete parts of the texture that were not vis-

	FID			SSIM			Feature- ℓ_1		
	Projection	NVS	Texture Field	Projection	NVS	Texture Field	Projection	NVS	Texture Field
airplanes	15.375	17.816	9.236	0.970	0.964	0.968	0.143	0.158	0.136
cars	70.070	72.209	24.271	0.840	0.887	0.885	0.236	0.238	0.192
chairs	8.045	8.788	5.791	0.931	0.947	0.941	0.127	0.125	0.124
tables	10.254	9.311	8.846	0.934	0.953	0.943	0.123	0.117	0.123
benches	16.527	17.484	12.965	0.921	0.936	0.937	0.144	0.142	0.133
cabinets	23.374	44.753	24.251	0.895	0.933	0.932	0.158	0.166	0.148
sofas	17.151	20.187	9.618	0.898	0.938	0.938	0.167	0.152	0.135
lamps	19.214	20.324	16.747	0.954	0.961	0.960	0.162	0.168	0.161
vessels	30.400	44.461	21.483	0.928	0.935	0.929	0.194	0.209	0.185
mean	23.379	28.370	14.801	0.919	0.939	0.937	0.162	0.164	0.149

Table 1: **Single Image Texture Reconstruction.** Quantitative Evaluation using the FID, SSIM and Feature- ℓ_1 metrics.

	FID	SSIM	Feature- ℓ_1
Projection	65.745	0.850	0.229
Im2Avatar	141.209	0.734	0.281
NVS	73.223	0.870	0.228
Texture Field	59.424	0.870	0.217

Table 2: **Full Pipeline.** Quantitative Evaluation using the FID, SSIM and Feature- ℓ_1 metrics. Results per category are presented in Section 2 of the supplementary material. Note that Im2Avatar was only applied to cars, chairs and tables.

ible in the input view. At the same time, our model successfully transfers texture regions visible in the input view to the shape (e.g. windows and tires of the car). In contrast, both the projection baseline and NVS show considerable artifacts. While the projection baseline transfers texture incorrectly into ambiguous regions, NVS often leads to unrealistic image artifacts such as rough edges and bleeding colors. A quantitative comparison is provided in Table 1. While NVS achieves the best SSIM, our method performs best in terms of FID and Features- ℓ_1 distance. This is consistent with the qualitative result in Fig. 4, as SSIM is a local score whereas FID and the Features- ℓ_1 distance are global scores that better capture visual similarity and realism.

Full Pipeline: In order to obtain a full single-view textured 3D reconstruction pipeline, we combine Texture Fields with Occupancy Networks [25]. For a fair comparison, we also combine the projection baseline and the NVS baseline with the output from [25].

Our qualitative results in Fig. 7 and Fig. 9 demonstrate that our approach is able to reconstruct 3D models with texture from just a single view of the model. In contrast to Im2Avatar⁷ [39], NVS and the projection baseline, we ob-

⁷Unfortunately, Sun et al. provide training data only for a subset of our training data. We therefore train and evaluate Im2Avatar on the training and test split provided by Sun et al. As the test split provided by Sun et al.

serve that our method achieves more consistent and realistic looking outputs. This is also reflected in the numerical results in Table 2: while NVS and our method both obtain the best SSIM, our method achieves the best FID and Features- ℓ_1 distances.

Real images: In Fig. 5, we finally investigate whether our approach is also able to transfer texture information from real input images to ground truth CAD models. Towards this goal, we apply our approach to the images provided by [38] and [49] and select CAD models that are similar to the object shown in the input image. We observe that our model generalizes reasonably well to real data despite only being trained on synthetic data.

4.3. Unconditional Model

In this section, we conduct unconditional experiments to investigate if Texture Fields can also be applied in a purely generative manner where we only provide the shape of the object to the network, but not the 2D image. Towards this goal, we train both, the VAE and the GAN model, on the “cars” category. During training, we provide both target images and depth maps but no input views to the network. During testing, we sample the latent code z from a standard normal distribution to obtain random texture samples for the given 3D object.

Random samples for the GAN and the VAE model are shown in Fig. 8. While our unconditional models successfully generate realistic textures, both the VAE and GAN samples contain artifacts similar to those present when applying VAEs and GANs to the image domain. For example, the samples of the VAE model are globally consistent but slightly blurry. In contrast, the samples for the GAN model are sharper but contain artifacts. In the future, we would like to explore combinations of VAEs and GANs and more

and our test split are disjoint for the “chairs” and “tables” categories, we can therefore also only show a qualitative comparison on “cars” category.



Figure 6: **Latent Space Interpolations (VAE)**. Our generative model learns a meaningful latent space embedding.

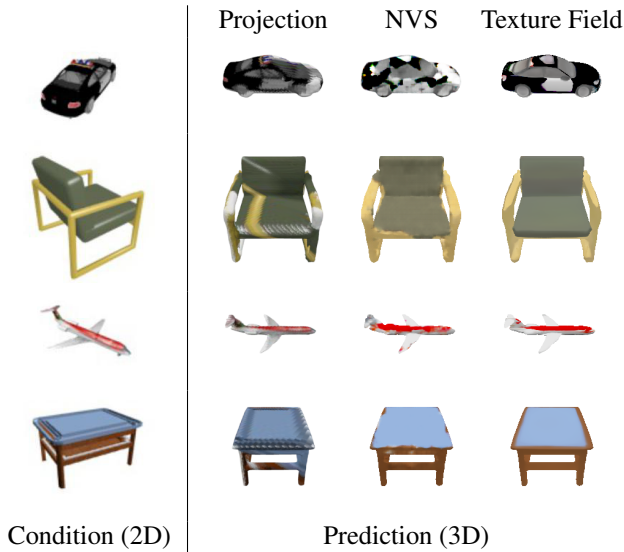


Figure 7: **Full Pipeline**. Results using Occupancy Networks [25] for 3D reconstruction in combination with projection, NVS and Texture Fields for appearance estimation.




Figure 8: **Generative Model**. Textures generated using the GAN (top 2 rows) and VAE (bottom 2 rows) models. Note that no 2D image is provided as input in this setting.

advanced models and training methods [12, 20] to improve upon these initial results.

Fig. 6 shows interpolations in the latent space for the VAE model. We see that the VAE has learned a meaningful latent space and can hence interpolate smoothly between different texture samples. Moreover, in Fig. 10, we demon-

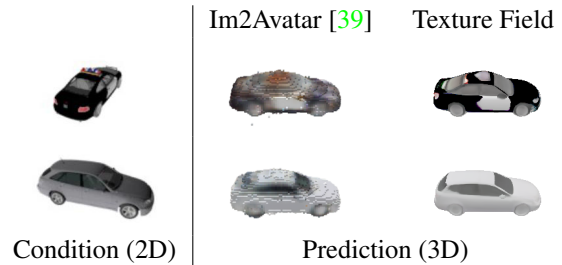


Figure 9: **Voxels vs. Function Space**. We compare our full pipeline against a voxel-based reconstruction method [39].

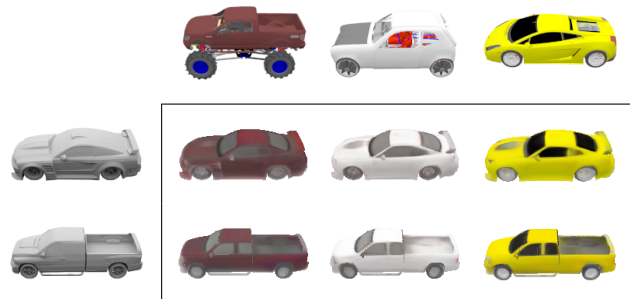


Figure 10: **Texture Transfer (VAE)**. Our model transfers appearance information (top) to the target models (left).

strate that our VAE model also allows for successfully transferring texture from one model to another one.

5. Conclusion

In this paper we introduced Texture Fields, a novel continuous representation for texture of 3D shapes. Our experiments show that Texture Fields can predict high frequency textures from just a single object view. Moreover, we have demonstrated that Texture Fields can also be used in an unconditional setting where we are only given the shape of the 3D object. We hence believe that Texture Fields are a useful representation for 3D reconstruction and hope that they will become an integral part of the next generation of 3D generative models.

Acknowledgements

This work was supported by Microsoft Research through its PhD Scholarship Program, by the Intel Network on Intelligent Systems and by a NVIDIA research gift.

References

- [1] Hassan Alhaija, Siva Mustikovela, Andreas Geiger, and Carsten Rother. Geometric image synthesis. In *Proc. of the Asian Conf. on Computer Vision (ACCV)*, 2018. 2
- [2] Hassan Alhaija, Siva Mustikovela, Lars Mescheder, Andreas Geiger, and Carsten Rother. Augmented reality meets computer vision: Efficient data generation for urban driving scenes. *International Journal of Computer Vision (IJCV)*, 126(9):961–972, 2018. 2
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2019. 2, 5
- [4] André Brock, Theodore Lim, James M. Ritchie, and Nick Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv.org*, 1608.04236, 2016. 1
- [5] Angel X. Chang, Thomas A. Funkhouser, Leonidas J. Guibas, Pat Hanrahan, Qi-Xing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiang Xiao, Li Yi, and Fisher Yu. Shapenet: An information-rich 3d model repository. *arXiv.org*, 1512.03012, 2015. 5
- [6] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2
- [7] Christopher Bongsoo Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2016. 1, 2, 5
- [8] Jiankang Deng, Shiyang Cheng, Niannan Xue, Yuxiang Zhou, and Stefanos Zafeiriou. UV-GAN: adversarial facial UV map completion for pose-invariant face recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [9] Haoqiang Fan, Hao Su, and Leonidas J. Guibas. A point set generation network for 3d object reconstruction from a single image. *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [10] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J. Rezende, S. M. Ali Eslami, and Yee Whye Teh. Neural processes. In *Proc. of the International Conf. on Machine Learning (ICML) Workshops*, 2018. 2
- [11] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NIPS)*, 2014. 2, 3, 4
- [12] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. DRAW: A recurrent neural network for image generation. In *Proc. of the International Conf. on Machine Learning (ICML)*, 2015. 4, 8
- [13] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. AtlasNet: A papier-mâché approach to learning 3d surface generation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4, 5
- [15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 5
- [16] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *Proc. of the International Conf. on Machine Learning (ICML)*, 2017. 5
- [17] Zeng Huang, Tianye Li, Weikai Chen, Yajie Zhao, Jun Xing, Chloe LeGendre, Linjie Luo, Chongyang Ma, and Hao Li. Deep volumetric video from very sparse multi-view performance capture. In *Proc. of the European Conf. on Computer Vision (ECCV)*, pages 351–369, 2018. 2
- [18] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2018. 1, 2
- [19] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2018. 2, 4, 5
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 5, 8
- [21] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2015. 5
- [22] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *Proc. of the International Conf. on Learning Representations (ICLR)*, 2014. 3, 4, 5
- [23] Yiyi Liao, Simon Donne, and Andreas Geiger. Deep marching cubes: Learning explicit surface representations. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [24] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *Proc. of the International Conf. on Machine Learning (ICML)*, 2018. 2, 4, 5
- [25] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [26] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv.org*, 1411.1784, 2014. 4
- [27] Eunbyung Park, Jimei Yang, Ersin Yumer, Duygu Ceylan, and Alexander C Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2

- [28] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [29] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 4, 5
- [30] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2016. 2, 4
- [31] Konstantinos Rematas, Chuong H. Nguyen, Tobias Ritschel, Mario Fritz, and Tinne Tuytelaars. Novel views of objects from a single image. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(8):1576–1590, 2017. 2
- [32] Danilo Jimenez Rezende, S. M. Ali Eslami, Shakir Mohamed, Peter Battaglia, Max Jaderberg, and Nicolas Heess. Unsupervised learning of 3d structure from images. In *Advances in Neural Information Processing Systems (NIPS)*, 2016. 1
- [33] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proc. of the International Conf. on Machine Learning (ICML)*, 2014. 4, 5
- [34] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015. 5
- [35] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. *arXiv.org*, 1905.05172, 2019. 2
- [36] Shunsuke Saito, Lingyu Wei, Liwen Hu, Koki Nagano, and Hao Li. Photorealistic facial texture inference using deep neural networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [37] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in Neural Information Processing Systems (NIPS)*, 2015. 4
- [38] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B Tenenbaum, and William T Freeman. Pix3d: Dataset and methods for single-image 3d shape modeling. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 7
- [39] Yongbin Sun, Ziwei Liu, Yue Wang, and Sanjay E. Sarma. Im2avatar: Colorful 3d reconstruction from a single image. *arXiv.org*, 1804.06375, 2018. 1, 2, 5, 7, 8
- [40] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Multi-view 3d models from single images with a convolutional network. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2016. 2
- [41] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 2012. 5
- [42] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [43] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2Mesh: Generating 3D mesh models from single RGB images. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2018. 1, 2
- [44] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. on Image Processing (TIP)*, 13(4):600–612, 2004. 5
- [45] Jiajun Wu, Yifan Wang, Tianfan Xue, Xingyuan Sun, Bill Freeman, and Josh Tenenbaum. MarrNet: 3D shape reconstruction via 2.5D sketches. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 1, 2
- [46] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in Neural Information Processing Systems (NIPS)*, 2016. 1
- [47] Wenqi Xian, Patsorn Sangkloy, Varun Agrawal, Amit Raj, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. TextureGAN: Controlling deep image synthesis with texture patches. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [48] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2016. 2
- [49] Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Joshua B. Tenenbaum, and William T. Freeman. Visual object networks: Image generation with disentangled 3D representations. *Advances in Neural Information Processing Systems (NIPS)*, 2018. 2, 7