

---

# Equivariant Neural Rendering

---

Emilien Dupont<sup>1</sup> Miguel Angel Bautista<sup>2</sup> Alex Colburn<sup>2</sup> Aditya Sankar<sup>2</sup> Joshua Susskind<sup>2</sup> Qi Shan<sup>2</sup>

## Abstract

We propose a framework for learning neural scene representations directly from images, without 3D supervision. Our key insight is that 3D structure can be imposed by ensuring that the learned representation *transforms* like a real 3D scene. Specifically, we introduce a loss which enforces equivariance of the scene representation with respect to 3D transformations. Our formulation allows us to infer and render scenes in real time while achieving comparable results to models requiring minutes for inference. In addition, we introduce two challenging new datasets for scene representation and neural rendering, including scenes with complex lighting and backgrounds. Through experiments, we show that our model achieves compelling results on these datasets as well as on standard ShapeNet benchmarks.

## 1. Introduction

Designing useful 3D scene representations for neural networks is a challenging task. While several works have used traditional 3D representations such as voxel grids (Maturana & Scherer, 2015; Nguyen-Phuoc et al., 2018; Zhu et al., 2018), meshes (Jack et al., 2018), point clouds (Qi et al., 2017; Insafutdinov & Dosovitskiy, 2018) and signed distance functions (Park et al., 2019), they each have limitations. For example, it is often difficult to scalably incorporate texture, lighting and background into these representations. Recently, neural scene representations have been proposed to overcome these problems (Eslami et al., 2018; Sitzmann et al., 2019a;b), usually by incorporating ideas from graphics rendering into the model architecture.

In this paper, we argue that equivariance with respect to 3D transformations provides a strong inductive bias for neural rendering and scene representations. Indeed, we argue that,

---

<sup>1</sup>University of Oxford, UK <sup>2</sup>Apple Inc, USA. Correspondence to: Emilien Dupont <dupont@stats.ox.ac.uk>, Qi Shan <qshan@apple.com>.




Figure 1. From a single image (left), our model infers a scene representation and generates new views of the scene (right) with a learned neural renderer.

for many tasks, scene representations need not be explicit (such as point clouds and meshes) as long as they *transform* like explicit representations.

Our model is trained with no 3D supervision and only requires images and their relative poses to learn equivariant scene representations. Our formulation does not pose any restrictions on the rendering process and, as result, we are able to model complex visual effects such as reflections and cluttered backgrounds. Unlike most other scene representation models (Eslami et al., 2018; Sitzmann et al., 2019a;b), our model does not require any pose information at inference time. From a single image, we can infer a scene representation, transform it and render it (see Fig. 1). Further, we can infer and render scene representations in real time while many scene representation algorithms require minutes to perform inference from an image or a set of images (Nguyen-Phuoc et al., 2018; Sitzmann et al., 2019b; Park et al., 2019).

While several works achieve impressive results by training models on images of a single scene and then generating novel views of that same scene (Mildenhall et al., 2020), we focus on generalizing across different scenes. This provides an additional challenge as we are required to learn a prior over shapes and textures to generalize to novel scenes. Our approach also allows us to bypass the need for different scenes in the training set to be aligned (or share the same coordinate system). Indeed, since we learn scene representations that *transform* like real scenes, we only require *relative* transformations to train the model. This is particularly advantageous when considering real scenes with complicated backgrounds where alignment can be difficult to achieve.

Neural rendering and scene representation models are usually tested and benchmarked on the ShapeNet dataset (Chang et al., 2015). However, the images produced from this dataset are often very different from real scenes: they are rendered on empty backgrounds and only involve a single fixed object. As our model does not rely on 3D supervision, we are able to train it on rich data where it is very expensive or difficult to obtain 3D ground truths. We therefore introduce two new datasets of posed images which can be used to test models with complex visual effects. The first dataset, MugsHQ, is composed of photorealistic renders of colored mugs on a table with an ambient background. The second dataset, 3D mountains, contains renders of more than 500 mountains in the Alps using satellite and topography data. In summary, our contributions are:

- We introduce a framework for learning scene representations and novel view synthesis without explicit 3D supervision, by enforcing equivariance between the change in viewpoint and change in the latent representation of a scene.
- We show that we can generate convincing novel views in real time without requiring alignment between scenes nor pose at inference time.
- We release two new challenging datasets to test representations and neural rendering for complex, natural scenes, and show compelling rendering results on each, highlighting the versatility of our method.

## 2. Related Work

**Scene representations.** Traditional scene representations (e.g. point clouds, voxel grids and meshes) do not scale well due to memory and compute requirements. Truncated signed distance functions (SDF) have been used to aggregate depth measurements from 3D sensors (Curless & Levoy, 1996) to map and track surfaces in real-time (Newcombe et al., 2011), without requiring assumptions about surface structure. Nießner et al. extend these implicit surface methods by incrementally fusing depth measurements into a hashed memory structure. More recently Park et al. extend SDF representations to whole classes of shapes, with learned neural mappings. Similar implicit neural representations have been used for 3D reconstruction from a single view (Xu et al., 2019a; Mescheder et al., 2018).

**Neural rendering.** Neural rendering approaches produce photorealistic renderings given noisy or incomplete 3D or 2D observations. In Thies et al., incomplete 3D inputs are converted to rich scene representations using neural textures, which fill in and regularize noisy measurements. Sitzmann et al. encode geometry and appearance into a latent code that

is decoded using a differentiable ray marching algorithm. Similar to our work, DeepVoxels (Sitzmann et al., 2019a) encodes scenes into a 3D latent representation. In contrast with our work, these methods either require 3D information during training, complicated rendering priors or expensive inference schemes.

**Novel view synthesis.** In Eslami et al., one or more input views with camera poses are aggregated into a context feature vector, and are rendered into a target 2D image given a query camera pose. Tobin et al. extend this base method using epipolar geometrical constraints to improve the decoding. Our model does not require the expensive sequential decoding steps of these models and enforces 3D structure through equivariance. Tatarchenko et al. can perform novel view synthesis for single objects consistent with a training set, but require depth to train the model. Hedman et al.; Hedman et al.; Thies et al.; Xu et al. use coarse geometric proxies. Our method only requires images and their poses to train, and can therefore extend more readily to real scenes with minimal assumptions about geometry. Works based on flow estimation for view synthesis (Sun et al., 2018; Zhou et al., 2016) predict a flow field over the input image(s) conditioned on a camera viewpoint transformation. These approaches model a free-form deformation in image space, as a result, they cannot explicitly enforce equivariance with respect to 3D rotation. In addition, these models are commonly restricted to single objects, not entire scenes.

**Equivariance.** While translational equivariance is a natural property of convolution on the spatial grid, traditional neural networks are not equivariant with respect to general transformation groups. Equivariance for discrete rotations can be achieved by replicating and rotating filters (Cohen & Welling, 2016a). Equivariance to rotation has been extended to 3D using spherical CNNs (Esteves et al., 2017). Steerable filters (Cohen & Welling, 2016b) and equivariant capsule networks (Lenssen et al., 2018) achieve approximate smooth equivariance by estimating pose and transforming filters, or by disentangling pose and filter representations. Worrall et al. use equivariance to learn autoencoders with interpretable transformations, although they do not explicitly encode 3D structure in the latent space. Olszewski et al.’s method is closely related to ours but only focuses on a limited range of transformations, instead of complete 3D rotations. In our method, we achieve equivariance by treating our latent representation as a geometric 3D data structure and applying rotations directly to this representation.

## 3. Equivariant Scene Representations

We denote an image by  $\mathbf{x} \in \mathcal{X} = \mathbb{R}^{c \times h \times w}$  where  $c, h, w$  are the number of channels, height and width of the image respectively. We denote a scene representation by  $\mathbf{z} \in \mathcal{Z}$ . We further define a rendering function  $g : \mathcal{Z} \rightarrow \mathcal{X}$  mapping




Figure 2. Left: A camera on the sphere observing an explicit scene representation (a mesh). Right: A camera on the sphere observing an implicit scene representation (a 3D tensor).

scene representations to images and an inverse renderer  $f : \mathcal{X} \rightarrow \mathcal{Z}$  mapping images to scenes.

We distinguish between two classes of scene representations: explicit and implicit representations (see Fig. 2). Explicit representations are designed to be interpreted by humans and are rendered by a fixed interpretable process. As an example,  $\mathbf{z}$  can be a 3D mesh and  $g$  a standard rendering function such as a raytracer. Implicit representations, in contrast, are abstract and need not be human interpretable. For example,  $\mathbf{z}$  could be the latent space of an autoencoder and  $g$  a neural network. We argue that, for many tasks, scene representations need not be explicit as long as they *transform* like explicit representations.

Indeed, we can consider applying some transformation  $T^{\mathcal{Z}}$  to a scene representation. For example, we can rotate and translate a 3D mesh. The resulting image rendered by  $g$  should then reflect these transformations, that is we would expect an equivalent transformation  $T^{\mathcal{X}}$  to occur in image space (see Fig. 3). We can write down this relation as

$$T^{\mathcal{X}} g(\mathbf{z}) = g(T^{\mathcal{Z}} \mathbf{z}). \quad (1)$$

This equation encodes the fact that transforming a scene representation with  $T^{\mathcal{Z}}$  and rendering it with  $g$  is equivalent to rendering the original scene and performing a transformation  $T^{\mathcal{X}}$  on the rendered image. More specifically, the renderer is equivariant with respect to the transformations in image and scene space<sup>1</sup>. We then define an *equivariant scene representation* as one that satisfies the equivariance relation in equation (1). We can therefore think of equivariant scene representations as a generalization of several other scene representations. Indeed, meshes, voxels, point clouds (and so on) paired with their appropriate rendering function all satisfy this equation.

## 4. Model

In this section, we design a model and loss that can be used to learn equivariant scene representations from data.

<sup>1</sup>Formally,  $T^{\mathcal{X}}$  and  $T^{\mathcal{Z}}$  represent the action of a group, such as the group of 3D rotations  $\text{SO}(3)$  or the group of 3D rotations and translations  $\text{SE}(3)$ .




Figure 3. Rotating a mesh with  $T^{\mathcal{Z}}$  and rendering it with  $g$  is equivalent to rendering the original mesh and applying a transformation  $T^{\mathcal{X}}$  in image space. This is true regardless of the choice of scene representation and rendering function.

While our formulation applies to general transformations and scene representations, we focus on the case where the scene representations are deep voxels and the family of transformations is 3D rotations. Specifically, we set  $\mathcal{Z} = \mathbb{R}^{c_s \times d_s \times h_s \times w_s}$  where  $c_s$ ,  $d_s$ ,  $h_s$ ,  $w_s$  are the channels, depth, height and width of the scene representation. We denote the rotation operation in scene space by  $R^{\mathcal{Z}}$  and the equivalent rotation operation acting on rendered images  $\mathbf{x}$  by  $R^{\mathcal{X}}$ .

As our model learns implicit scene representations, we do not require 3D ground truths. Instead, our dataset is composed of pairs of views of scenes and relative camera transformations linking the two views. Specifically, we assume the camera observing the scenes is on a sphere looking at the origin. For a given scene, we consider two image captures of the scene  $\mathbf{x}_1$  and  $\mathbf{x}_2$  and the *relative* camera transformation between the two  $\theta = \theta \hat{\mathbf{n}}$  where  $\theta$  is the angle and  $\hat{\mathbf{n}}$  the axis parameterizing the 3D rotation<sup>2</sup>. A training data point is then given by  $(\mathbf{x}_1, \mathbf{x}_2, \theta)$ . In practice, we capture a large number of views for each scene and randomly sample new pairs at every iteration in training. This allows us to build models that generalize well across a large variety of camera transformations.

To design a loss that enforces equivariance with respect to the rotation transformation, we consider two images of the *same scene* and their relative transformation  $(\mathbf{x}_1, \mathbf{x}_2, \theta)$ . We first map the images through the inverse renderer to obtain their scene representations  $\mathbf{z}_1 = f(\mathbf{x}_1)$  and  $\mathbf{z}_2 = f(\mathbf{x}_2)$ . We then rotate each encoded representation by its relative transformation  $R_{\theta}^{\mathcal{Z}}$ , such that  $\tilde{\mathbf{z}}_1 = R_{\theta}^{\mathcal{Z}} \mathbf{z}_1$  and  $\tilde{\mathbf{z}}_2 = (R_{\theta}^{\mathcal{Z}})^{-1} \mathbf{z}_2$ . As  $\mathbf{z}_1$  and  $\mathbf{z}_2$  represent the same scene in different poses, we expect the rotated  $\tilde{\mathbf{z}}_1$  to be rendered

<sup>2</sup>We use the axis-angle parameterization for notational convenience, but any rotation formalism such as euler angles, rotation matrices and quaternions could be used. In our implementation, we parameterize this rotation by a rotation matrix.




Figure 4. Model training. We encode two images  $\mathbf{x}_1, \mathbf{x}_2$  of the same scene into their respective scene representations  $\mathbf{z}_1, \mathbf{z}_2$ . Since they are representations of the same scene viewed from different points, we can rotate each one into the other. The rotated scene representations  $\tilde{\mathbf{z}}_1, \tilde{\mathbf{z}}_2$  should then be decoded to match the swapped image pairs  $\mathbf{x}_2, \mathbf{x}_1$ .

as the image  $\mathbf{x}_2$  and the rotated  $\tilde{\mathbf{z}}_2$  as  $\mathbf{x}_1$ . This is illustrated in Fig. 4. We can then ensure our model obeys these transformations by minimizing

$$\mathcal{L}_{\text{render}} = \|\mathbf{x}_2 - g(\tilde{\mathbf{z}}_1)\| + \|\mathbf{x}_1 - g(\tilde{\mathbf{z}}_2)\|. \quad (2)$$

As  $\mathbf{x}_2 = R_{\theta}^{\mathbf{x}} \mathbf{x}_1$ , minimizing this loss then corresponds to satisfying the equivariance property for the renderer  $g$ . Note that the form of this loss function is similar to the ones proposed by Worrall et al. and Olszewski et al.

**Model architecture.** In contrast to most other works learning implicit scene representations (Worrall et al., 2017; Es-lami et al., 2018; Chen et al., 2019), our representation is spatial in three dimensions, allowing us to use fully convolutional architectures for both the inverse and forward neural renderer. To build the forward renderer, we take inspiration from RenderNet (Nguyen-Phuoc et al., 2018) and HoloGAN (Nguyen-Phuoc et al., 2019) as these have been shown to achieve good performance on rendering tasks. Specifically, the scene representation  $\mathbf{z}$  is mapped through a set of 3D convolutions, followed by a projection layer of  $1 \times 1$  convolutions and finally a set of 2D convolutions mapping the projection to an image. The inverse renderer is simply defined as the transpose of this architecture (see Fig. 5). For complete details of the architecture, please refer to the appendix.

**Voxel rotation.** Defining the rotation operation in scene space  $R^{\mathcal{Z}}$  is crucial. As our scene representation  $\mathbf{z}$  is a deep voxel grid, we simply apply a 3D rotation matrix to the coordinates of the features in the voxel grid. As the rotated points may not align with the grid, we use inverse warping with trilinear interpolation to reconstruct the values at the voxel locations (see Szeliski for more detail). We note that warping and interpolation operations are available




Figure 5. Model architecture. An input image (top left) is mapped through 2D convolutions (blue), followed by an inverse projection (purple) and a set of 3D convolutions (green). The inferred scene is then rendered through the transpose of this architecture.

in frameworks such as Pytorch and Tensorflow, making it simple to implement voxel rotations in practice.

**Rendering loss.** There are several possible choices for the rendering loss, the most common being the  $\ell_1$  norm,  $\ell_2$  norm and SSIM (Wang et al., 2004) or combinations thereof. As noted in other works (Worrall et al., 2017; Snell et al., 2017) a weighted sum of  $\ell_1$  and SSIM works well in practice. However, we found that our model is not particularly sensitive to the choice of regression loss, and analyse the various trade offs through ablation studies in the experimental section.

## 5. Experiments

We perform experiments on ShapeNet benchmarks (Chang et al., 2015) as well as on two new datasets designed to challenge the model on more complex scenes. For all experiments, the images are of size  $128 \times 128$  and the scene representations are of size  $64 \times 32 \times 32 \times 32$ . For both the 2D and 3D parts of the network we use residual layers for convolutions that preserve the dimension of the input and strided convolutions for downsampling layers. We use the LeakyReLU nonlinearity (Maas et al.) and GroupNorm (Wu & He, 2018) for normalization. Complete architecture and training details can be found in the appendix.

Most novel view synthesis works are tested on the ShapeNet dataset or variants of it. However, renders from ShapeNet objects are typically very far from real life scenes, which tends to limit the use cases for models trained on them. As our scene representation and rendering framework make no restricting assumptions about the rendering process (such as requiring single objects, no reflections, no background etc.), we create new datasets to test the performance of our model on more advanced tasks.

The new datasets are challenging by design and are com-

	TCO	dGQN	SRN	OURS
REQUIRES ABSOLUTE POSE	YES	YES	YES	NO
REQUIRES POSE AT INFERENCE TIME	NO	YES	YES	NO
OPTIMIZATION AT INFERENCE TIME	NO	NO	YES	NO

Table 1. Requirements for each baseline. Our model performs comparably to other models that make much stronger assumptions about the data and inference process.

posed of photorealistic 3D scenes and 3D landscapes with textures from satellite images. We achieve compelling results on these datasets and hope they will spur further research into scene representations that are not limited to simple scenes without backgrounds. The code and datasets are available at <https://github.com/apple/ml-equivariant-neural-rendering>.

### 5.1. Baselines

We compare our model with three strong baselines. The first is the model proposed by Tatarchenko et al. which we refer to as TCO, the second is a deterministic variant of Generative Query Networks (Eslami et al., 2018) which we refer to as dGQN and the third is the Scene Representation Network (SRN) as described in Sitzmann et al.<sup>3</sup>. All baselines make strong assumptions that substantially simplify the view-synthesis and scene representation problem. We discuss each of these assumptions in detail below and provide a comparison in Table 1. Our model requires neither of these assumptions, making the task it has to solve considerably more challenging while also being more generally applicable.

**Absolute and relative pose.** All baselines require an absolute coordinate system<sup>4</sup> for the pose (or viewpoints). For example, when trained on chairs, the viewpoint corresponding to the camera being at the origin would be the one observing the chair face on. The poses are then absolute in the sense that the camera at the origin corresponds to observing the chair face on for all chairs, i.e. we need all scenes to be perfectly aligned. While this is possible for simple datasets like ShapeNet, it is difficult to define a consistent alignment for a set of scenes, particularly for complex scenes with backgrounds and real life images. In contrast, our model does not require any notion of alignment or absolute pose. Equivariance is exactly why we are able to build a representation that is “origin-free”, because it only depends on relative transformations between poses.

<sup>3</sup>For detailed descriptions of these baselines, please refer to the appendix of (Sitzmann et al., 2019b).

<sup>4</sup>This is often referred to as a world coordinate system.



Figure 6. Novel view synthesis for chairs. Given a single image of an unseen object (left), we infer a scene representation, rotate and render it with our learned renderer to generate novel views. Due to space constraints we include chairs with interesting properties here and show randomly sampled chairs in the appendix.

**Pose at inference time.** In order to infer a scene representation, our model takes as input a single image of the scene. In contrast, both dGQN and SRNs require an image as well as the viewpoint from which the image was taken. This considerably simplifies the task as the model does not need to infer the pose.

**Optimization at inference time.** At inference time, SRNs require solving an optimization problem in order to fit a scene to the model. As such, inferring a scene representation from a single input image (on a Tesla V100 GPU) takes 2 minutes with SRNs but only 22ms for our model (three orders of magnitude faster). The idea of training at inference time is a crucial element of SRNs and other works in 3D computer vision (Park et al., 2019), but is not required for our model.

### 5.2. Chairs

We evaluate our model on the ShapeNet chairs class by following the experimental setup given in Sitzmann et al., using the same train/validation/test splits. The dataset is composed of 6591 chairs each with 50 views sampled uniformly on the sphere for a total of 329,550 images. Images are sampled on the *full sphere* around the object, making the task much more difficult than typical setups which limit the elevation or azimuth or both (Tatarchenko et al., 2016; Chen et al., 2019; Olszewski et al., 2019).

**Novel view synthesis.** Results for novel view synthesis are shown in Fig. 6. The novel views were produced by taking a single image of an *unseen chair*, inferring its scene representation with the inverse renderer, rotating the scene and generating a novel view with the learned neural renderer. As can be seen, our model is able to generate plausible views of new chairs even when viewed from difficult angles and in the presence of occlusion. The model works well even for oddly shaped chairs with thin structures.

DATASET	TCO	dGQN	SRN	OURS
CHAIRS	21.27	21.59	22.89	22.83

Table 2. Reconstruction accuracy (higher is better) in PSNR (units of dB) for baselines and our model on ShapeNet chairs.

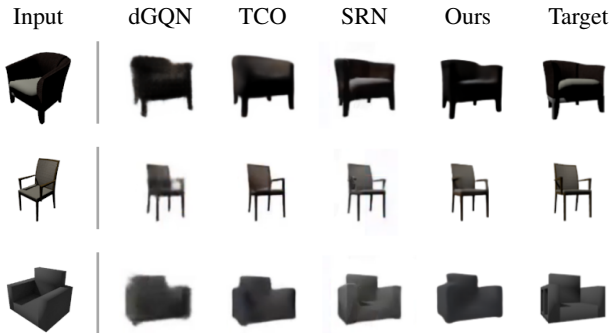


Figure 7. Qualitative comparisons for single shot novel view synthesis. The baseline images were borrowed with permission from Sitzmann et al.

**Quantitative comparisons.** To perform quantitative comparisons, we follow the setup in Sitzmann et al. by considering a single informative view of an unseen test object and measuring the reconstruction performance on the upper hemisphere around the object (results are shown in Table 2). Surprisingly, even though our model makes much weaker assumptions than all the baselines, it significantly improves upon both the TCO and dGQN baselines and is comparable with the state of the art SRNs.

**Qualitative comparisons.** We show qualitative comparisons with the baselines for single shot novel view synthesis in Fig. 7. As can be seen our model produces high quality novel views that are comparable to or better than dGQN and TCO while being slightly worse than SRNs.

### 5.3. Cars

We also evaluate our model on the ShapeNet cars class, allowing us to test our model on images with richer texture than chairs. The dataset is composed of 3514 cars each with 50 views sampled uniformly on the sphere for a total of 175,700 images.

**Novel view synthesis.** As can be seen in Fig. 8, our model is able to generate plausible views for cars with various colors and thin structures like spoilers. While our model successfully infers 3D shape and appearance, it still struggles to capture some fine texture and geometry details (see Section 6 for a thorough discussion of the limitations and failures of our model).

**Absolute and relative poses.** As mentioned in Section 5.1,



Figure 8. Novel view synthesis for cars.



Figure 9. Qualitative comparisons on cars between SRNs, SRNs with relative poses around the up axis and our model.

our model only relies on relative transformations and therefore alleviates the need for alignment between scenes. As all baselines require absolute poses and alignment between scenes, we run tests to see how important this assumption is. Specifically, we break the alignment between scenes in the cars dataset by randomly rotating each scene around the up axis<sup>5</sup>. We then train an SRN model on the perturbed and unperturbed dataset to understand to which extent the model relies on the absolute coordinates. As can be seen in Fig. 9, breaking the alignment between scenes significantly deteriorates the performance of SRNs while it leaves the performance of our model unchanged. This is similarly reflected when measuring reconstruction accuracy on the test set (see Table 3).

### 5.4. MugsHQ

As the model does not make any restricting assumptions about the rendering process, we test it on more difficult scenes by building the MugsHQ dataset based on the mugs class from ShapeNet. Instead of rendering images on a blank background, every scene is rendered with an environment

<sup>5</sup>We found that rotating around one axis was enough to see a significant effect. Rotating around all 3 axes would likely have an even larger effect.

DATASET	SRN	SRN (RELATIVE)	OURS
CARS	22.36	21.05	22.26

Table 3. Reconstruction accuracy (higher is better) in PSNR (units of dB) on ShapeNet cars.



Figure 10. Novel view synthesis on MugsHQ.

map (lighting conditions) and a checkerboard disk platform. For each of the 214 mugs, we sample 150 viewpoints uniformly over the upper hemisphere and render views using the Mitsuba renderer (Jakob, 2010). Note that the environment map and disk platform is the same for every mug. The resulting scenes include more complex visual effects like reflections and look more realistic than typical ShapeNet renders, making the task of novel view synthesis considerably more challenging. A complete description of the dataset as well as samples can be found in the appendix.

**Novel view synthesis.** Results for single shot novel view synthesis on unseen mugs are shown in Fig. 10. As can be seen, the model successfully infers the shape of unseen mugs from a single image and is able to perform large viewpoint transformations. Even from difficult viewpoints, the model is able to produce consistent and realistic views of the scenes, even generating reflections on the mug edges. As is the case for the ShapeNet dataset, our model can still miss fine details such as thin mug handles and struggles with some oddly shaped mugs (see Section 6 for examples).

### 5.5. Mountains

We also introduce 3D mountains, a dataset of mountain landscapes. We created the dataset by scraping the height, latitude and longitude of the 559 highest mountains in the Alps (we chose this mountain range because it was easiest to find data). We then used satellite images combined with topography data to sample random views of each mountain at a fixed height (see appendix for samples and detailed description). This dataset is extremely challenging, with varied and complex geometry and texture. While obtaining high quality results on this dataset is beyond the scope of




Figure 11. Novel view synthesis on 3D mountains.



Figure 12. Comparisons on chairs showing the trade off between different rendering losses.

our algorithm, we hope it can be useful for pushing the boundaries of research in neural rendering.

**Novel view synthesis.** Results for single shot novel view synthesis are shown in Fig. 11. While the model struggles to capture high frequency detail, it faithfully reproduces the 3D structure and texture of the mountain as the camera rotates around the scene representation. For a variety of mountain landscapes (snowy, rocky etc.), our model is able to generate plausible, albeit blurry, views. An interesting feature is that, for views near the input image, the generated images are considerably sharper than for views far away from the input. This is likely due to the considerable uncertainty in generating views far from the source view: given the front of a mountain, there are many plausible ways the back of the mountain could appear. As our model is deterministic, it generates sharper views near the input where there is less uncertainty and blurs views far from the input where there is more uncertainty.

### 5.6. Ablation studies

We perform ablation studies to test the trade offs between various rendering losses. Fig. 12 shows the difference in generated images when using  $\ell_2$  and  $\ell_1 + \text{SSIM}$  losses. While both losses perform well, the  $\ell_2$  loss produces somewhat blurrier images than the  $\ell_1 + \text{SSIM}$  loss. However, there are also cases where the  $\ell_1 + \text{SSIM}$  produces artifacts that the  $\ell_2$  loss does not. Ultimately, there is a trade

off between using the two losses and the choice is largely dependent on the application.

## 6. Scope, limitations and future work

In this section, we discuss some of the advantages and weaknesses of our method as well as potential directions for future work.

**Advantages.** The main advantage of our model is that it makes very few assumptions about the scene representation and rendering process. Indeed, we learn representations simply by enforcing equivariance with respect to 3D rotations. As such, we can easily encode material, texture and lighting which is difficult with traditional 3D representations. The simplicity of our model also means that it can be trained purely from posed 2D images with no 3D supervision. As we have shown, this allows us to apply our method to interesting data where obtaining 3D geometry is difficult. Crucially, and unlike most other methods, our model does not require alignment between scenes nor any pose information at inference time. Further, our model is fast: inferring a scene representation simply corresponds to performing a forward pass of a neural network. This is in contrast to most other methods that require solving an expensive optimization problem at inference time for every new observed image (Nguyen-Phuoc et al., 2018; Park et al., 2019; Sitzmann et al., 2019b). Rendering is also performed in a single forward pass, making it faster than other methods that often require recurrence to produce an image (Eslami et al., 2018; Sitzmann et al., 2019b).

**Limitations.** As our scene representation is spatial and 3-dimensional, our model is quite memory hungry. This implies we need to use a fairly small batch size which can make training slow (see appendix for detailed analysis of training times). Using a voxel-like representation could also make it difficult to generalize the model to other symmetries such as translations. In addition, our model typically produces samples of lower quality than models which make stronger assumptions. As an example, SRNs generally produce sharper and more detailed images than our model and are able to infer more fine-grained 3D information. Further SRNs can, unlike our model, generalise to viewpoints that were not observed during training (such as rolling the camera or zooming). While this is partly because we are solving a task that is inherently more difficult, it would still be desirable to narrow this gap in performance. We also show some failure cases of our model in Fig. 13. As can be seen, the model struggles with very thin structures as well as objects with unusual shapes. Further, the model can create unrealistic renderings in certain cases, such as mugs with disconnected handles.

**Future work.** The main idea of the paper is that equivariance with respect to symmetries of a real scene provides




Figure 13. Failure examples of our model. As can be seen, the model fails on oddly shaped chairs, cars and mugs. On cars, the model sometimes infers the correct shape but misses high frequency texture detail. On mugs, the model can miss mug handles and other thin structure.

a strong inductive bias for representation learning of 3D environments. While we implement this using voxels as the representation and rotations as the symmetry, we could just as well have chosen point clouds as the representation and translation as the symmetry. The formulation of the model and loss are independent of the specific choices of representation and symmetry and we plan to explore the use of different representations and symmetries in future work.

In addition, our model is deterministic, while inferring a scene from an image is an inherently uncertain process. Indeed, for a given image, there are several plausible scenes that could have generated it and, similarly, several different scenes could be rendered as the same image. It would therefore be interesting to learn a *distribution* over scenes  $p(\text{scene}|\text{image})$ . Training a probabilistic or adversarial model may also help sharpen rendered images.

Another promising route would be to use the learned scene representation for 3D reconstruction. Indeed, most 3D reconstruction methods are object-centric (i.e. every object is reconstructed in the same orientation). This has been shown to cause models to effectively perform shape classification instead of reconstruction (Tatarchenko et al., 2019). As our scene representation is view-centric, it is likely that it could be useful for the downstream task of 3D reconstruction in the view-centric case.

## 7. Conclusion

In this paper, we proposed learning scene representations by ensuring that they *transform* like real 3D scenes. The proposed model requires no 3D supervision and can be trained using only posed 2D images. At test time, our model can, from a single image and in real time, infer a scene representation and manipulate this representation to render novel views. Finally, we introduced two challenging new datasets which we hope will help spur further research into neural rendering and scene representations for complex scenes.



## Acknowledgements

We thank Shuangfei Zhai, Walter Talbott and Leon Gatys for useful discussions. We also thank Lilian Liang and Leon Gatys for help with running compute jobs. We thank Per Fahlberg for his help in generating the 3D mountains dataset. We are also grateful for Vincent Sitzmann for his help with generating ShapeNet datasets and benchmarks. We also thank Russ Webb for feedback on an early version of the manuscript. Finally we thank the anonymous reviewers for their useful feedback and suggestions.

## References

- Chang, A. X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Chen, X., Song, J., and Hilliges, O. Monocular neural image based rendering with continuous view control. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- Cohen, T. and Welling, M. Group equivariant convolutional networks. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2990–2999, New York, New York, USA, 20–22 Jun 2016a. PMLR.
- Cohen, T. S. and Welling, M. Steerable cnns, 2016b.
- Curless, B. and Levoy, M. A volumetric method for building complex models from range images. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, pp. 303–312, New York, NY, USA, 1996. Association for Computing Machinery. ISBN 0897917464. doi: 10.1145/237170.237269. URL <https://doi.org/10.1145/237170.237269>.
- Eslami, S. M. A., Jimenez Rezende, D., Besse, F., Viola, F., Morcos, A. S., Garnelo, M., Ruderman, A., Rusu, A. A., Danihelka, I., Gregor, K., Reichert, D. P., Buesing, L., Weber, T., Vinyals, O., Rosenbaum, D., Rabinowitz, N., King, H., Hillier, C., Botvinick, M., Wierstra, D., Kavukcuoglu, K., and Hassabis, D. Neural scene representation and rendering. *Science*, 360(6394), 2018.
- Esteves, C., Allen-Blanchette, C., Makadia, A., and Daniilidis, K. Learning  $so(3)$  equivariant representations with spherical cnns. *CoRR*, 2017. URL <http://arxiv.org/abs/1711.06721>.
- Hedman, P., Ritschel, T., Drettakis, G., and Brostow, G. Scalable inside-out image-based rendering. *ACM Transactions on Graphics*, 35, 2016.
- Hedman, P., Philip, J., Price, T., Frahm, J.-M., Drettakis, G., and Brostow, G. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics*, 37, 2018.
- Insafutdinov, E. and Dosovitskiy, A. Unsupervised learning of shape and pose with differentiable point clouds. In *Advances in neural information processing systems*, pp. 2802–2812, 2018.
- Jack, D., Pontes, J. K., Sridharan, S., Fookes, C., Shirazi, S., Maire, F., and Eriksson, A. Learning free-form deformations for 3d object reconstruction. In *Asian Conference on Computer Vision*, pp. 317–333. Springer, 2018.
- Jakob, W. Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>.
- Lenssen, J. E., Fey, M., and Libuschewski, P. Group equivariant capsule networks. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 8844–8853. Curran Associates, Inc., 2018.
- Maas, A. L., Hannun, A. Y., and Ng, A. Y. Rectifier nonlinearities improve neural network acoustic models.
- Maturana, D. and Scherer, S. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 922–928. IEEE, 2015.
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., and Geiger, A. Occupancy networks: Learning 3d reconstruction in function space, 2018.
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. Nerf: Representing scenes as neural radiance fields for view synthesis. *arXiv preprint arXiv:2003.08934*, 2020.
- Newcombe, R. A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A. J., Kohli, P., Shotton, J., Hodges, S., and Fitzgibbon, A. W. Kinectfusion: Real-time dense surface mapping and tracking. In *ISMAR*, volume 11, pp. 127–136, 2011.
- Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., and Yang, Y.-L. Hologan: Unsupervised learning of 3d representations from natural images. *arXiv preprint arXiv:1904.01326*, 2019.
- Nguyen-Phuoc, T. H., Li, C., Balaban, S., and Yang, Y. Rendernet: A deep convolutional network for differentiable rendering from 3d shapes. In *Advances in Neural Information Processing Systems*, pp. 7891–7901, 2018.

- Nießner, M., Zollhöfer, M., Izadi, S., and Stamminger, M. Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (ToG)*, 32(6):169, 2013.
- Olszewski, K., Tulyakov, S., Woodford, O., Li, H., and Luo, L. Transformable bottleneck networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 7648–7657, 2019.
- Park, J. J., Florence, P., Straub, J., Newcombe, R. A., and Lovegrove, S. DeepSDF: Learning continuous signed distance functions for shape representation. *CoRR*, 2019.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 652–660, 2017.
- Sitzmann, V., Thies, J., Heide, F., Niessner, M., Wetzstein, G., and Zollhofer, M. Deepvoxels: Learning persistent 3d feature embeddings. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019a.
- Sitzmann, V., Zollhöfer, M., and Wetzstein, G. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *arXiv preprint arXiv:1906.01618*, 2019b.
- Snell, J., Ridgeway, K., Liao, R., Roads, B. D., Mozer, M. C., and Zemel, R. S. Learning to generate images with perceptual similarity metrics. In *2017 IEEE International Conference on Image Processing (ICIP)*, pp. 4277–4281. IEEE, 2017.
- Sun, S.-H., Huh, M., Liao, Y.-H., Zhang, N., and Lim, J. J. Multi-view to novel view: Synthesizing novel views with self-learned confidence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 155–171, 2018.
- Szeliski, R. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
- Tatarchenko, M., Dosovitskiy, A., and Brox, T. Multi-view 3d models from single images with a convolutional network. In *European Conference on Computer Vision*, pp. 322–337. Springer, 2016.
- Tatarchenko, M., Richter, S. R., Ranftl, R., Li, Z., Koltun, V., and Brox, T. What do single-view 3d reconstruction networks learn? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3405–3414, 2019.
- Thies, J., Zollhöfer, M., Theobalt, C., Stamminger, M., and Nießner, M. Ignor: Image-guided neural object rendering. *arXiv*, 2018.
- Thies, J., Zollhöfer, M., and Nießner, M. Deferred neural rendering: Image synthesis using neural textures. *SIGGRAPH*, 2019.
- Tobin, J., Zaremba, W., and Abbeel, P. Geometry-aware neural rendering. In *Advances in Neural Information Processing Systems 32*, pp. 11555–11565. Curran Associates, Inc., 2019.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- Worrall, D. E., Garbin, S. J., Turmukhambetov, D., and Brostow, G. J. Interpretable transformations with encoder-decoder networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5726–5735, 2017.
- Wu, Y. and He, K. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 3–19, 2018.
- Xu, Q., Wang, W., Ceylan, D., Mech, R., and Neumann, U. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction, 2019a.
- Xu, Z., Bi, S., Sunkavalli, K., Hadap, S., Su, H., and Ramamoorthi, R. Deep view synthesis from sparse photometric images. *ACM Transactions on Graphics*, 38, 2019b.
- Zhou, T., Tulsiani, S., Sun, W., Malik, J., and Efros, A. A. View synthesis by appearance flow. In *European conference on computer vision*, pp. 286–301. Springer, 2016.
- Zhu, J.-Y., Zhang, Z., Zhang, C., Wu, J., Torralba, A., Tenenbaum, J., and Freeman, B. Visual object networks: Image generation with disentangled 3d representations. In *Advances in neural information processing systems*, pp. 118–129, 2018.