

Control-NeRF: Editable Feature Volumes for Scene Rendering and Manipulation

Verica Lazova¹²

verica.lazova@uni-tuebingen.de

Vladimir Guzov¹²

vladimir.guzov@mnf.uni-tuebingen.de

Kyle Olszewski³

kolszewski@snap.com

Sergey Tulyakov³

stulyakov@snap.com

Gerard Pons-Moll¹²

gerard.pons-moll@uni-tuebingen.de


¹University of Tübingen²Max Planck Institute for Informatics, Saarland Informatics Campus³Snap Inc.

Figure 1: Given a set of posed input images, our model allows for novel view synthesis (a) and controlled scene editing. For example, inserting objects from one scene into another (b); or multiplying, moving or deleting objects within the scene (c)

Abstract

We present *Control-NeRF*¹, a method for performing flexible, 3D-aware image content manipulation while enabling high-quality novel view synthesis, from a set of posed input images. NeRF-based approaches [23] are effective for novel view synthesis, however such models memorize the radiance for every point in a scene within a neural network. Since these models are scene-specific and lack a 3D scene representation, classical editing such as shape manipulation, or combining scenes is not possible. While there are some recent hybrid approaches that combine NeRF with external scene representations such as sparse voxels, planes, hash tables, etc. [16, 5, 24, 9], they focus mostly on efficiency and don't explore the scene editing and manipulation capabilities of hybrid approaches. With the aim of exploring controllable scene representations for novel view synthesis, our model couples learnt scene-specific 3D feature

volumes with a general NeRF rendering network. We can generalize to novel scenes by optimizing only the scene-specific 3D feature volume, while keeping the parameters of the rendering network fixed. Since the feature volumes are independent of the rendering model, we can manipulate and combine scenes by editing their corresponding feature volumes. The edited volume can then be plugged into the rendering model to synthesize high-quality novel views. We demonstrate scene manipulations including: scene mixing; applying rigid and non-rigid transformations; inserting, moving and deleting objects in a scene; while producing photo-realistic novel-view synthesis results.

1. Introduction

Scene manipulation and rendering are long-standing problems in graphics, with the goal of creating the desired visual content and providing immersive abilities to explore it. Traditionally, the process consists of acquiring

¹<https://virtualhumans.mpi-inf.de/control-nerf/>

textured meshes of objects and scenes, followed by combining them using specialized software and hardware to reach the desired composition, and finally, rendering the scene using graphical pipelines. Acquiring, composing and rendering are non-trivial problems that require time and experience, and, hence, are not readily available for amateur users. Impressive progress in Novel View Synthesis (NVS) sparked by the recently introduced Neural Radiance Fields (NeRF) [23], represents an attractive vehicle for scene manipulation and rendering. However, NeRF and most follow-up works suffer from two main shortcomings, limiting their use for creative applications. First, they require per-scene training, and second, the scene is represented by a neural network, which makes *editing* and manipulation difficult. Recent work has shown how to generalize NeRF to novel scenes [54], but those works have not demonstrated editing and control. Other recent work have shown editing capabilities by learning per object NeRF models or decomposing a single scene into foreground objects and background [39, 58, 17]. However, these models are either object or scene specific, or work on synthetic scenes without realistic background [16], limiting its applicability. What is missing is a neural representation and model which allows to represent multiple *real* scenes, while allowing intuitive control. This would retain the realism and simplicity of neural rendering models, while keeping the versatility and intuitive control of traditional computer graphics representations (meshes, volumes and textures). In this work, we present Control-NeRF, a novel approach which can represent multiple scenes, and allows intuitive control and editing. We learn a latent representation of the scene, encoded as a spatially disentangled feature volume (*i.e.*, in which the point features describe the content and radiance at that point in the scene), coupled with a neural rendering function that computes the radiance and density conditioned on the point feature. This decouples the rendering network from the neural scene representation which results in several advantages. First, the model can be trained on multiple scenes at once, producing different scene representations for each of them, while learning a general rendering network. Second, once the model is learned, new scene representations can be learned while holding the rendering network fixed (desirable if, for example, we want to stream scenes without having to re-train or transmit the rendering network). Third, as we show in the experiments, the learned representations are aligned with the real 3D scenes, which allows for intuitive manipulation such as displacing, rotating and transforming objects, integrating objects from other scenes, or simply combining scenes, see Fig. 1. Most importantly, because each scene has its own representation and the rendering network is shared across scenes, editing and composition can be done post-hoc without re-training. We demonstrate that learning Control-NeRF efficiently on real scenes

requires a careful coarse-to-fine strategy — in which the optimized feature volume dimensions are progressively increased — and a total variation regularizer on the feature volume representation. In our evaluations, we demonstrate that our approach allows for NVS using a single model for multiple real training scenes while being comparable to scene specific models. We also demonstrate how to efficiently generalize to novel scenes by optimizing the scene representation while keeping the rendering network fixed. Finally, we demonstrate various creative manipulation tasks such as compositing of different real scenes, displacing and rotating objects, and inserting objects. In summary, our primary contributions are:

- We demonstrate that the learned scene representations are aligned with the real 3D scenes allowing for easy, simple and creative 3D-aware editing without having to retrain the model.
- We show that techniques such as multi-resolution training and total variation regularization are essential for efficiently optimizing the 3D volume containing the scene.
- Extensive evaluations, demonstrating our method significantly outperforms competing approaches in terms of the types of manipulations possible and quality of results.
- We will release our code and trained models for research purposes.

2. Related Work

Novel View Synthesis. Novel-view-synthesis (NVS) is a widely studied problem in the area of image-based rendering. Most NVS methods are focused on warping or blending the input images and inpainting the occluded regions. Recent NVS efforts [33, 34, 1, 10, 31] have achieved high quality results relying on geometry proxies, such as rough reconstructions, depth maps or point-clouds to warp the input images to the target view. Many works use the ability of generative networks to hallucinate occluded regions from one or a few images [48, 41, 21], which can be complemented with the use of appearance flow [64, 27, 45]. However, in case of large viewpoint transformations, best results are attained only for simple/synthetic scenes, or by using a large number of input images. Other techniques, *e.g.* multiplane images [63, 57], have proven suitable for large-scale scenes like those captured in real photographs. Most of these frameworks, however, focus only on NVS and typically provide little or no ability to edit the scene content (*e.g.*, adding or deforming objects).

Implicit Surface and Appearance Representations. The use of implicit surfaces [20, 28, 4, 6] for geometry

and appearance reconstruction has proven popular in recent works, with their ability to capture detailed objects with varying topology at arbitrary resolutions. Methods such as PIFu [36] use these to capture the surface and texture of dynamic humans from monocular images, an approach that was refined and improved in [37, 15, 11] to allow for higher fidelity and realtime performance capture. Methods like these can be used for NVS simply by rendering the obtained reconstructions. However, while they achieve impressive results for individual objects, they struggle to capture the full geometry and appearance of complex real scenes.

Volumetric Representations. Explicit voxel grids [13] have recently been employed for various tasks related to implicit surface and appearance representation, including generative modeling of 3D objects [51, 65], shape and appearance reconstruction from images [12, 52, 46, 26, 14]. Other recent works, such as [40] have explored the use of latent representations with a volumetric structure to implicitly encode a scene’s appearance and structure for neural rendering. They use multiple images of static objects to learn a feature volume that can be resampled to a given camera’s viewpoint. Also, [18] use multiple calibrated images of static and dynamic scenes to learn a latent volumetric representation that can be used for rendering of novel views, including time varying effects (*e.g.* human motion). However, these works do not allow for interactive editing or manipulation of the scene, and are typically scene specific, requiring separate networks for each captured scene. Some methods, such as [26] train a network which infers a latent volumetric representation of previously unseen images that can be spatially transformed to allow for NVS and editing. However, the image quality of the manipulated objects is relatively low, and it only works on simple scenes.

Hybrid Latent and Geometric Representations. Other recent approaches combine explicit representations of a scene’s geometry with a latent representation to exploit neural rendering techniques. Some methods learn neural textures [49] used in conjunction with UV-maps to allow for realistic image synthesis and manipulation. The learnt textures, however, are specific to the corresponding objects and scenes used during training, and thus cannot generalize to new scenes without retraining. In NPBG [1], given several images of a scene with a corresponding 3D point cloud, neural descriptors are fitted to points, which are then used with the input data to learn to infer novel views of the scene. This work requires a point cloud of the scene, obtained using multi-view stereo or depth sensor data as part of the training process. Therefore the overall quality of the final results depends heavily on the quality of the reconstruction.

Neural Radiance Fields. Neural Radiance Fields (NeRF) [23] builds on prior work on implicit surface representations by introducing a sophisticated MLP architecture trained to produce an estimate of the density and out-

going radiance throughout the scene. Volume rendering techniques are used to enforce consistency with the training images, which enables the inference of high-quality novel views of the scene. Subsequent efforts have addressed various limitations of this work and extended it to new applications, *e.g.* accelerating its training and rendering performance and quality [47, 16, 32, 25, 54, 9, 44]; extending it to large-scale scenes [19, 61]; allowing for the capture and synthesis of dynamic scenes with non-rigid regions, including human heads and bodies [29, 30, 55]; relighting the captured content [2, 3, 42]; camera and body pose estimation [43, 59]; and NVS with unknown camera parameters [56]. Some works [50, 60, 7] use projected features from images into a space that may be queried in a manner similar to [36, 15]. With a NeRF-like radiance function they demonstrate the ability to perform NVS using a single or few input images. However the overall quality and complexity of the synthesized images is limited, and they do not enable general manipulations of the scene, as in our method. Recently there have been few works that combine voxel grids and neural radiance fields. [16] use sparse voxel fields to learn local radiance fields for improved rendering performance. For the given scene they build the voxels by pruning the voxel grid at training time. They can also do local shape editing and build scenes by compositing separate objects together. While this method shows impressive results on individual objects, they struggle to deal with real scenes with complex background and front facing scenes, where the scene is not observed from all sides. Another similar work, [58] has introduced a method that learns an object-compositional neural radiance field. They learn separately a scene branch to encode the scene appearance and individual object branches for all the object in the scene. This method allows for object-level editing, such as moving and transforming the objects in the scene. However unlike our method it is scene-specific and does not support moving objects across multiple scenes. For a more comprehensive survey of work in this area, please refer to [8].

3. Control-NeRF

We present our novel-view synthesis method, Control-NeRF, (Figure 2) that is based on feature volumes and Neural Radiance Fields [23] and allows for scene editing, mixing and manipulation. We decouple geometry/appearance from rendering by learning dense feature volume as representation for every scene and a single rendering model that generalizes across scenes. The rendering model takes a feature vector sampled from the volume and predicts density and color value. As shown in [23], these predictions are used as input to a volume rendering function that accumulates the point along a ray to generate a pixel color.

This section of the paper is organized as follows. In Sec. 3.1, we briefly review the general framework used

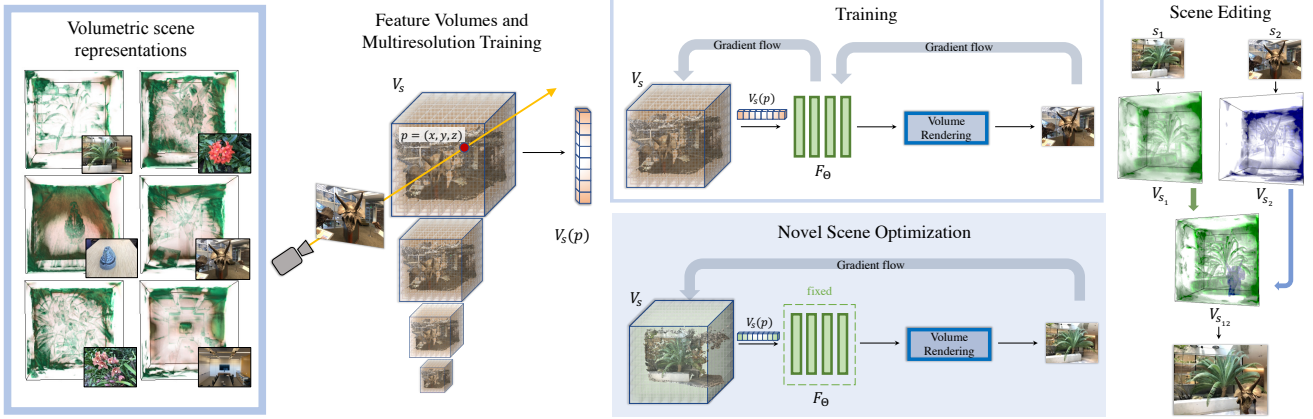


Figure 2: Our method learns a volumetric representations for multiple scenes simultaneously. Left in the figure we show visualizations of the learned feature volumes. We query the volume along the ray and predict color and density based on the obtained features. The pixel color is derived using volume rendering, similar to [23]. At training time the volume and the rendering network are trained jointly. For novel scenes, the rendering network is fixed and only the scene volume is optimized. As shown on the right, these volumes can be edited and mixed and for the purpose of scene editing.

for performing novel view synthesis using neural radiance fields. Sec. 3.2 describes how we make use of learnt feature volumes to condition the radiance field output on a given scene. In Sec. 3.3, we describe the losses and the training procedure for optimizing the network parameters and per-scene feature volumes. We also show how to learn feature volumes for novel scenes not seen at training time. (Sec. 3.3.4). Sec. 3.4 describes how we can use the learnt feature volumes for arbitrary creative scene manipulations and render the result.

3.1. Background

Most works based on Neural Radiance Fields [23] predict radiance and color for a pair of point and viewing angle direction of a single scene:

$$F_{\Theta} : (\gamma(\mathbf{p}), \gamma(\mathbf{d})) \rightarrow (\mathbf{c}, \sigma) \quad (1)$$

where $\mathbf{c} \in \mathbb{R}^3$ is an RGB value indicating the radiance from point $\mathbf{p} \in \mathbb{R}^3$ in direction $\mathbf{d} \in \mathbb{R}^3$, and $\sigma \in \mathbb{R}$ is the density value \mathbf{p} , indicating how much the radiance contributes to view rays intersecting the scene at that point. Optionally, one can use γ which is a positional encoding [53] used to allow this network to better capture high-frequency details. Images are rendered one pixel at a time, using volumetric sampling techniques, querying the MLP at points along the camera ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ (where \mathbf{o} indicates the camera origin and t indicates the distance from the origin along the ray) corresponding to that pixel. By integrating the radiance values at a point using its density, the appropriate color values can be computed.

The problem with NeRF based approaches is that the scene is memorized within the neural network, which makes

compositing of scenes and editing hard.

3.2. Formulation

To allow realistic editing, our method decouples the scene representation from the neural rendering network. Instead of memorizing a mapping from scene point and viewing directions to radiance with an MLP as in Eq. 1, we learn a scene-specific volume of deep features. Then a rendering network maps from deep features extracted at continuous locations of the volume, to radiance and color.

Scene Representation Given a set of input RGB images $\mathcal{I} = \{I_s^i\}_{i=1}^{N_s}$ from M training scenes $s \in \mathcal{S}$, $M = |\mathcal{S}|$, we seek to learn a latent volumetric representation $V_s \in \mathbb{R}^{WHD \times F}$ for each scene s , with a spatial resolution of $W \times H \times D$ and a feature vector of length F in each cell, which can be both rendered from novel views and edited to allow for novel manipulations of the scene content while still allowing for high-quality view synthesis. We use a resolution of $W = H = D = 128$ and a feature vector of length $F = 64$ in our experiments.

Rendering Network The rendering network is a learned mapping from a deep feature $\mathbf{v}_s \in \mathbb{R}^{64}$ to radiance and color. The deep feature describes the local shape and appearance of the corresponding position $\mathbf{p} = (x, y, z)$ in scene s extracted from a scene-specific volume of deep features. Mathematically,

$$F_{\Theta} : (S(V_s, \mathbf{p}), \gamma(\mathbf{d})) \rightarrow (\mathbf{c}, \sigma) \quad (2)$$

where the feature vector \mathbf{v}_s is obtained by sampling of the feature volume $\mathbf{v}_s = S(V_s, \mathbf{p})$, where S indicates the trilinear

ear resampling operation. As in NeRF the density F_Θ is integrated along rays \mathbf{r} to produce pixel colors – this operation is denoted by $\hat{C}(\mathbf{r}, V_s, \Theta)$. In contrast to equation 1, the formulation in equation 2 allows us to optimize the volume V_s for each scene, while simultaneously learning the parameters of the density network $F_\Theta : (\mathbf{v}_s, \gamma(\mathbf{d})) \rightarrow (\mathbf{c}, \sigma)$. After this initial training stage, the parameters Θ of this rendering module are fixed. For every novel scene, we only optimize its feature volume V_s . This will allow us to combine and edit scenes by manipulating their respective feature volumes V_s , and render the result using the general rendering module \hat{C} .

Note: As in NeRF, in practice, 2 networks are trained: a coarse network in which samples are taken from evenly-spaced intervals along the view ray, and a fine network which uses the density values from the coarse network to select sample points more likely to contribute to the corresponding ground-truth pixel color value $C(\mathbf{r})$. In the following, we denote the density integrals of the coarse and fine networks as $\hat{C}_f(\mathbf{r})$, $\hat{C}_c(\mathbf{r})$ respectively. ²

3.3. Training Strategy and Generalization

3.3.1 Training Losses

Reconstruction Loss. Our primary loss is a straightforward reconstruction loss on the rendered pixel values. As in [23], at each iteration we randomly sample and integrate a subset of the rays \mathcal{R}_s from the images for the current scene s , and compute the mean-squared error between them and the corresponding pixels in the ground-truth images:

$$\mathcal{L}_r(\mathcal{R}_s, \mathcal{I}_s, V_s, \Theta) = \mathbb{E}_{\mathbf{r} \sim \mathcal{R}_s} \left[\left\| \hat{C}_c(\mathbf{r}, V_s, \Theta) - C(\mathbf{r}) \right\|_2^2 + \left\| \hat{C}_f(\mathbf{r}, V_s, \Theta) - C(\mathbf{r}) \right\|_2^2 \right] \quad (3)$$

Using this loss for each training scene (see Sec. 3.3.3), we jointly optimize the network parameters Θ and the feature volumes \mathcal{V} for all training scenes.

Total Variation Loss. One very useful property that we want the volumes to exhibit is local smoothness – neighbouring feature vectors should have similar values. NeRF [23], has this property by default, since it relies on \mathbb{R}^3 (3D locations as input). In order to encourage similar behaviour for our feature volumes we add regularization. In our experiments, we found that a more consistent and coherent feature volume was learned if we introduce a total variation regularization [35] loss to the learned feature volume. To reduce memory usage and computation, we apply this loss on the 64-dimensional feature vectors in a randomly sampled contiguous subregion $R \subset V_s$ that is 1/4 of the

current latent feature volume (see Sec. 3.3.2) for the current scene s during each training iteration.

$$\mathcal{L}_{tv}(V_s) = \mathbb{E}_{R \sim V_s} [|T(R)|], \quad (4)$$

$$T(R) = \sum_{i,j,k} \sqrt{|R_{i+1,j,k} - R_{i,j,k}|^2 + |R_{i,j+1,k} - R_{i,j,k}|^2 + |R_{i,j,k+1} - R_{i,j,k}|^2} \quad (5)$$

Thus, we minimize the following total loss function by optimizing the parameters Θ and feature volumes \mathcal{V} corresponding to the calibrated images \mathcal{I} and the corresponding view rays \mathcal{R} for each training scene:

$$\arg \min_{\mathcal{V}, \Theta} \mathcal{L}(\mathcal{R}, \mathcal{I}, \mathcal{V}, \Theta) = \mathbb{E}_{s \sim \mathcal{S}} [\mathcal{L}_r(\mathcal{R}_s, \mathcal{I}_s, V_s, \Theta) + \lambda \mathcal{L}_{tv}(V_s)] \quad (6)$$

where $\lambda = 10^{-4}$.

3.3.2 Multi-Resolution Volume Training

As the final volume contains a 64-dimensional feature vector per cell in the 128^3 volume, training the network at this full resolution is quite intensive. As such, we employ a hierarchical training process to compute these volumes in a coarse-to-fine manner. This allows for improved training time while retaining the ability to perform high-quality image synthesis and manipulation. We start training with a feature volume resolution of 16^3 . The model is trained until convergence, optimizing both the current feature volume V_s and rendering module parameters Θ . We then upsample the learnt feature volume to increase its dimensions by a factor of 2, and proceed to train until convergence at the new resolution. We use 4 stages in our hierarchical training process, doubling the feature volume dimensions at each stage until we reach the target resolution of 128^3 .

3.3.3 Multi-Scene Training

To allow the rendering module to be employed for multiple scenes, it needs to be trained in a multiscene scenario. During training we randomly select one of the scenes $s \in \mathcal{S}$ and load its feature volume V_s , then train using rays sampled from this volume for several consecutive iterations, before saving the feature volume and repeating the process with a new randomly selected scene. While sampling a new scene at each training iteration would better approximate the effect of incorporating samples from multiple scenes at each step in the optimization, this would require additional overhead as feature volumes are loaded into GPU memory, then copied back to be stored for their next use. We empirically found that 50 consecutive iterations between scene transitions produced a sufficient balance between training performance and multi-scene representation capacity.

²For simplicity, from here on we use Θ to refer to the total learnable parameters of the 2 rendering networks.

3.3.4 Generalization to Novel Scenes

After the initial training stage in which the parameters Θ of the radiance network F_Θ are trained in conjunction with the optimization of the M per-scene feature volumes V_1, \dots, V_M , we allow for efficient generalization to novel scenes by fixing the parameters Θ and solely optimizing the parameters of the feature volumes corresponding to these novel scenes.

Given a new set of scenes \mathcal{G} not used during the initial training stage, and a set of images corresponding to each scene $\mathcal{I}' = \{I_g^i\}_{i=1}^{N_g}$ for each scene $g \in \mathcal{G}$, we perform the optimization process as described above, while *only* optimizing the corresponding feature volume V_g' for each scene. We employ the hierarchical training strategy defined in Sec. 3.3.2, and the losses defined in Eqns. 3 and 4, but for these scenes only optimize the feature volumes corresponding to each scene g to minimize the total loss:

$$\arg \min_{V'} \mathcal{L}(\mathcal{R}', \mathcal{I}', V', \Theta) = \mathbb{E}_{g \sim \mathcal{G}} [\mathcal{L}_r(\mathcal{R}_g', \mathcal{I}_g', V_g', \Theta) + \lambda \mathcal{L}_{tv}(V_g')] \quad (7)$$

Given sufficient training scenes, the learnt radiance function can be applied to optimize for novel scenes more efficiently than when training to infer the volumes and network parameters together as in the initial training process. In our experiments we show that a small number of training scenes (only 6) are sufficient to train a generalizable radiance function.

3.4. Scene Editing and Manipulation

Our volumetric representation of scene-specific content allows for scene manipulations by editing its feature volume. We can swap features between different feature volumes for mixing scenes, copy features to make duplicates of an object or zero-out features to make deletions. Rigid and non-rigid transformations of an object can be applied by resampling the volume. By applying trilinear resampling to contiguous subregions of the feature volume (or the entire volume, if global scene deformations are desired), nonrigid spatial manipulations can be applied. If V_o is the original feature volume and $P \in \mathbb{R}^{3 \times W \times H \times D}$ is a matrix of 3D coordinates indicating where to sample from for each point in the modified volume, $V_m = S(V_o, P)$ will produce a volume with the desired spatial deformation, where S is the trilinear sampling operator.

4. Experiments and Results

4.1. Dataset and Implementation Details

For our initial training stage, we use 6 scenes from the dataset provided by LLFF [22], consisting of a total of 230 images (an average of approximately 38 images per scene) with the cameras' extrinsic and intrinsic parameters estimated by COLMAP [38]. After this stage, we fix the rendering module parameters and optimize the feature volumes for

	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NSVF [16]	20.414	0.536	0.449
NPBG [1]	19.430	0.727	0.242
Ours	25.635	0.853	0.181

Table 1: **Quantitative comparison with NPBG [1] and NSVF [16].** Metrics are computed across test images for scenes from LLFF [22] dataset. ‘‘Ours’’ is our method trained on 6 scenes simultaneously as in our original setup. Please consult the supplementary for more details.

new scenes individually. We use 2 scenes from this dataset, withheld during the initial training stage, to demonstrate our novel scene generalization capabilities (*fern* and *trex*, shown in Figure 1), consisting of a total of 75 images. Please consult the appendix for more details.

4.2. Scene Content Manipulation

Using the scene resampling and editing techniques described in Sec. 3.4, we demonstrate various creative manipulations enabled by our method. In Figure 5 we show scene manipulation by moving object from one scene into another. The scenes shown in these examples are real scenes from the LLFF [22] dataset. In Figure 4, we show single scene editing by removing objects or making copies of existing objects. Please keep in mind that all editing results are obtained only by mixing or shifting the features within the 3D volumes. No fine-tuning or post-processing steps are used, as we want to show the 3D editing capabilities of our method, without any adjustments in image space.

4.3. Evaluations and Comparisons.

Editing Comparisons. We provide qualitative and quantitative comparisons of our approach to scene manipulation to two related methods, Neural Point-Based Graphics (NPBG) [1] and Neural Sparse Voxel Fields (NSVF) [16]. NPBG uses 3D point clouds of a scene with corresponding RGB images and camera poses to allow for both realistic neural rendering of novel views of the scene and copying content from one scene into another. NSVF uses sparse voxels scene representation that is pruned at training time. This representation is useful for isolated objects, but struggles with real scenes with complex background.

For our quantitative comparisons we evaluate the novel-view-synthesis capabilities of our method in comparison to NPBG and NSVF (Table 1). For the qualitative comparisons we evaluate the scene editing capabilities of our method in comparison to NPBG (Figure 5). While in theory NSVF could perform similar manipulations, the official implementation doesn't support multi-scene editing or scene manipulation. Nevertheless in Figure 3 we compare our method to NSVF in the task of novel view synthesis of complex real scenes. We could not compare to the recent




Figure 3: **Comparison to NSVF [16]** in Novel view synthesis. As discussed, NSVF struggles with real frontal scenes, in which the content is not captured from 360°.

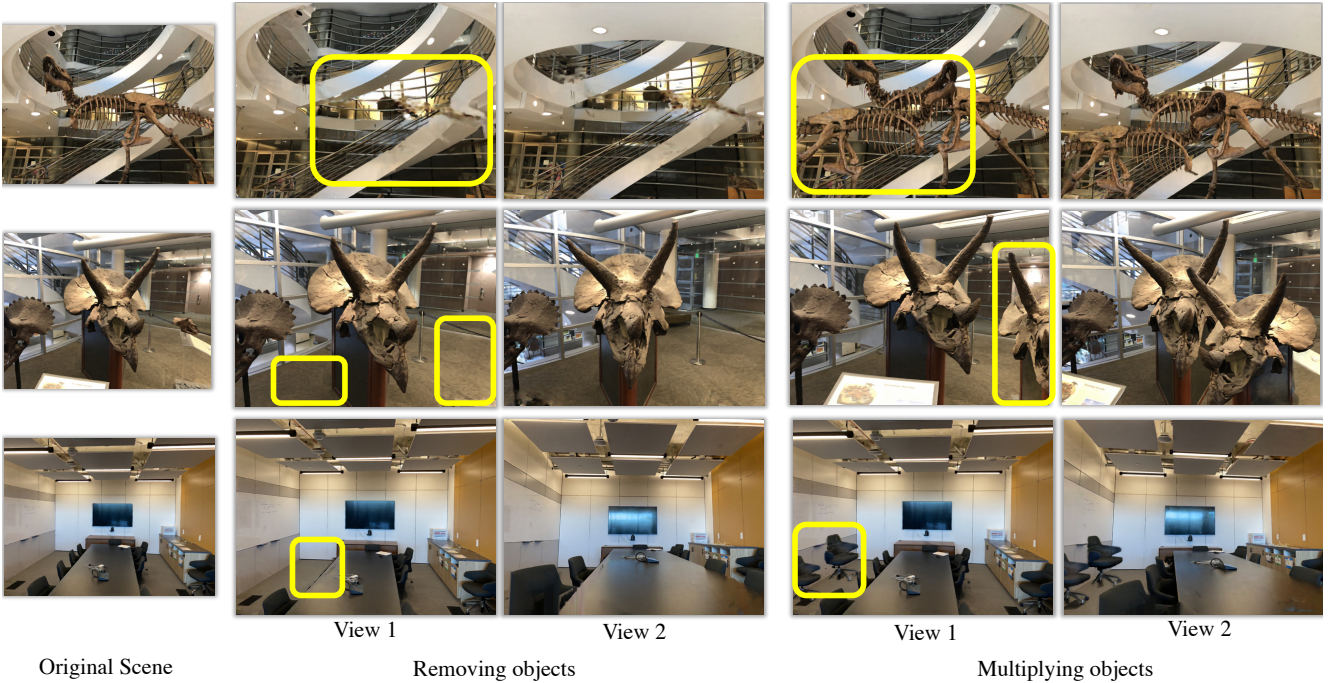


Figure 4: **Replicating and removing object from scenes** The first column shows the original scene. The rest of the columns show the edited scene from two different views. The differences are marked with yellow rectangles in the first view.

method of [58], as there is currently no released implementation. Using 43 images from the 8 aforementioned scenes (The 6 initial training scenes and the 2 scenes optimized with fixed renderer parameters) withheld during the training process, we compute the Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [62] between the ground truth and the images synthesized using both methods. Table 1 contains the results, which show that our method outperforms [1] and [16] using each metric.

In Fig. 5 we show that, while both approaches can be used to combine scenes, our approach outperforms [1] when

it comes to editing capabilities. Please consult the supplementary video and document for animated results from these experiments, as well as further results and details on our approach and evaluations. We include additional NVS and manipulation results for various datasets. In all these scenarios, we only train the feature volumes of the new scenes while keeping the rendering parameters fixed. We also provide the results of experiments with non-rigid scene manipulation and articulated animation; an ablation study evaluating the utility of our approaches key components; a perceptual study asking users to evaluate our approach compared to NPBG [1]; and comparisons with the origi-



Figure 5: **Scene Editing Results.** The first column shows the original, unedited scenes. An object from the source scene is replicated in the target scene. The following columns show novel view of our editing results and results obtained using NPBG [1]. Please zoom in on these images and consult the supplementary video for more results and animations.

nal NeRF [23], as well as a modified version of this work adapted to handle multiple scenes.

5. Conclusion

With Control-NeRF, we explore a promising direction for flexible 3D scene manipulation using neural radiance fields. In disentangling the scene representation and rendering, our approach enables practical techniques for efficient scene editing and high-fidelity image synthesis. We demonstrate a wide range of such edits, *e.g.* replicating and removing objects, applying rigid and non-rigid transformations, and mixing scenes. One limitation is that our editing method does not explicitly handle shadows and different lighting condition between scenes. Thus, if scenes with vastly different lighting are combined, the results might not appear as convincing, a challenge we intend to address in future work. We also intend to explore methods to enhance our approach with more editing capabilities: modifying the

textures and appearance of scene; changing and adapting to different lighting, adding shadows; exploring techniques for more efficient scene optimization and rendering.

Acknowledgements: We thank Aymen Mir, Bharat Bhatnagar, Garvita Tiwari, Ilya Petrov, Jan Eric Lenssen, Julian Chibane, Keyang Zhou and Xiaohan Zhang for the in-depth discussions, valuable insights and honest feedback.

This work is and supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A; and partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 409792180 (Emmy Noether Programme, project: Real Virtual Humans).

Gerard Pons-Moll is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645.

References

- [1] Kara-Ali Aliev, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. *arXiv preprint arXiv:1906.08240*, 2(3):4, 2019. 2, 3, 6, 7, 8
- [2] Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sulkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Neural reflectance fields for appearance acquisition. <https://arxiv.org/abs/2008.03824>, 2020. 3
- [3] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik Lensch. NeRD: Neural reflectance decomposition from image collections. <https://arxiv.org/abs/2012.03918>, 2020. 3
- [4] Rohan Chabra, Jan Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local SDF priors for detailed 3D reconstruction. In *The European Conference on Computer Vision (ECCV)*, 2020. 2
- [5] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. *arXiv preprint arXiv:2112.07945*, 2021. 1
- [6] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion, 2020. 2
- [7] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo radiance fields (srf): Learning view synthesis for sparse views of novel scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7911–7920, June 2021. 3
- [8] Frank Dellaert and Lin Yen-Chen. Neural volume rendering: Nerf and beyond, 2021. 3
- [9] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 1, 3
- [10] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (TOG)*, 37(6):1–15, 2018. 2
- [11] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. ARCH: Animatable Reconstruction of Clothed Humans. In *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3090–3099, Seattle, WA, USA, June 2020. IEEE. 3
- [12] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, and Andrew Fitzgibbon. Kinectfusion: Real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, UIST '11, page 559–568, New York, NY, USA, 2011. Association for Computing Machinery. 3
- [13] James T. Kajiya and Brian P. Von Herzen. Ray tracing volume densities. *Computer Graphics (SIGGRAPH)*, 1984. 3
- [14] Markus Knoche, István Sáradi, and Bastian Leibe. Reposing humans by warping 3D features. In *CVPR Workshop on Towards Human-Centric Image/Video Synthesis*, 2020. 3
- [15] Ruilong Li, Yuliang Xiu, Shunsuke Saito, Zeng Huang, Kyle Olszewski, and Hao Li. Monocular real-time volumetric performance capture. *arXiv preprint arXiv:2007.13988*, 2020. 3
- [16] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020. 1, 2, 3, 6, 7
- [17] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields, 2021. 2
- [18] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 2019. 3
- [19] Ricardo Martin-Brualla, Noha Radwan, Mehdi Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the wild: Neural radiance fields for unconstrained photo collections. <https://arxiv.org/abs/2008.02268>, 2020. 3
- [20] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy Networks: Learning 3D reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [21] Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rerendering in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6878–6887, 2019. 2
- [22] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019. 6
- [23] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020. 1, 2, 3, 4, 5, 8
- [24] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. 1
- [25] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Chakravarty R. Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. Donerf: Towards real-time rendering of neural radiance fields using depth oracle networks, 2021. 3
- [26] Kyle Olszewski, Sergey Tulyakov, Oliver Woodford, Hao Li, and Linjie Luo. Transformable bottleneck networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 3

- [27] Eunbyung Park, Jimei Yang, Ersin Yumer, Duygu Ceylan, and Alexander C Berg. Transformation-grounded image generation network for novel 3d view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [28] Jeong Joon Park, Pete Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019. 2
- [29] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan Goldman, Steven Seitz, and Ricardo Martin-Brualla. Deformable neural radiance fields. <https://arxiv.org/abs/2011.12948>, 2020. 3
- [30] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. <https://arxiv.org/abs/2011.13961>, 2020. 3
- [31] Ruslan Rakhimov, Andrei-Timotei Ardelean, Victor Lempitsky, and Evgeny Burnaev. Npbg++: Accelerating neural point-based graphics. *arXiv preprint arXiv:2203.13318*, 2022. 2
- [32] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. DeRF: Decomposed radiance fields. <https://arxiv.org/abs/2011.12490>, 2020. 3
- [33] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *European Conference on Computer Vision*, pages 623–640. Springer, 2020. 2
- [34] Gernot Riegler and Vladlen Koltun. Stable view synthesis. *arXiv preprint arXiv:2011.07233*, 2020. 2
- [35] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms, 1992. 5
- [36] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 3
- [37] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2020. 3
- [38] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 6
- [39] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [40] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2019. 3
- [41] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *arXiv preprint arXiv:1906.01618*, 2019. 2
- [42] Pratul Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. NeRV: Neural reflectance and visibility fields for relighting and view synthesis. <https://arxiv.org/abs/2012.03927>, 2020. 3
- [43] Shih-Yang Su, Frank Yu, Michael Zollhoefer, and Helge Rhodin. A-NeRF: Surface-free human 3d pose refinement via neural rendering. <https://arxiv.org/abs/2102.06199>, 2021. 3
- [44] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. 3
- [45] Shao-Hua Sun, Minyoung Huh, Yuan-Hong Liao, Ning Zhang, and Joseph J. Lim. Multi-view to novel view: Synthesizing novel views with self-learned confidence. In *Proceedings of the European Conference on Computer Vision*, 2018. 2
- [46] Yongbin Sun, Ziwei Liu, Yue Wang, and Sanjay E. Sarma. Im2avatar: Colorful 3d reconstruction from a single image, 2018. 3
- [47] Matthew Tancik, Ben Mildenhall, Terrence Wang, Divi Schmidt, Pratul Srinivasan, Jonathan T. Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations. <https://arxiv.org/abs/2012.02189>, 2020. 3
- [48] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Single-view to multi-view: Reconstructing unseen views with a convolutional network. *CoRR abs/1511.06702*, 1(2):2, 2015. 2
- [49] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics 2019 (TOG)*, 2019. 3
- [50] Alex Trevithick and Bo Yang. GRF: Learning a general radiance field for 3D scene representation and rendering. <https://arxiv.org/abs/2010.04595>, 2020. 3
- [51] Shubham Tulsiani, Hao Su, Leonidas J. Guibas, Alexei A. Efros, and Jitendra Malik. Learning shape abstractions by assembling volumetric primitives. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [52] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 3
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 4
- [54] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-

- Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering, 2021. 2, 3
- [55] Ziyang Wang, Timur Bagautdinov, Stephen Lombardi, Tomas Simon, Jason Saragih, Jessica Hodgins, and Michael Zollhöfer. Learning compositional radiance fields of dynamic human heads. <https://arxiv.org/abs/2012.09955>, 2020. 3
- [56] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. NeRF-: Neural radiance fields without known camera parameters. <https://arxiv.org/abs/2102.07064>, 2021. 3
- [57] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion, 2021. 2
- [58] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13779–13788, 2021. 2, 3, 7
- [59] Lin Yen-Chen, Pete Florence, Jonathan T. Barron, Alberto Rodriguez, Phillip Isola, and Tsung-Yi Lin. iNeRF: Inverting neural radiance fields for pose estimation. <https://arxiv.org/abs/2012.05877>, 2020. 3
- [60] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. <https://arxiv.org/abs/2012.02190>, 2020. 3
- [61] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. NeRF++: Analyzing and improving neural radiance fields. <https://arxiv.org/abs/2010.07492>, 2020. 3
- [62] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric, 2018. 7
- [63] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Transactions on Graphics*, 37(4):65:1–65:12, 2018. 2
- [64] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *Proceedings of the European Conference on Computer Vision*, 2016. 2
- [65] Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Josh Tenenbaum, and Bill Freeman. Visual object networks: image generation with disentangled 3d representations. In *Proceedings of the Neural Information Processing Systems Conference*, pages 118–129, 2018. 3