

NeRFDeformer: NeRF Transformation from a Single View via 3D Scene Flows

Zhenggang Tang,¹ Zhongzheng Ren,¹ Xiaoming Zhao,¹ Bowen Wen,² Jonathan Tremblay²
 Stan Birchfield,² Alexander Schwing¹

¹University of Illinois Urbana-Champaign: {ztang, zr5, xz23, aschwing}@illinois.edu

²NVIDIA: {bowenw, jtremblay, sbirchfield}@nvidia.com

Abstract

We present a method for automatically modifying a NeRF representation based on a single observation of a non-rigid transformed version of the original scene. Our method defines the transformation as a 3D flow, specifically as a weighted linear blending of rigid transformations of 3D anchor points that are defined on the surface of the scene. In order to identify anchor points, we introduce a novel correspondence algorithm that first matches RGB-based pairs, then leverages multi-view information and 3D reprojection to robustly filter false positives in two steps. We also introduce a new dataset for exploring the problem of modifying a NeRF scene through a single observation. Our dataset¹ contains 113 synthetic scenes leveraging 47 3D assets. We show that our proposed method outperforms NeRF editing methods as well as diffusion-based methods, and we also explore different methods for filtering correspondences.

1. Introduction

Transforming a neural radiance field (NeRF) based on a single RGBD image is an important problem. Consider the field of robotics as an example, where NeRFs are often used to represent complicated 3D scenes [17, 35, 42, 48, 55]. Notably, whenever the scene is modified, the robot has to recapture multiple views to re-train a new NeRF. This process discards important information from the original scene and is time consuming. We are hence interested in developing tools that allow a given NeRF scene to be transformed into a new scene observed via a single RGBD image (see Fig. 1). Concretely, we are interested in retrieving the transformed scene geometry and rendering the new scene from different perspectives.

NeRF editing is a natural approach for solving this problem, and current works have shown tremendous success at modifying NeRF appearance [1] or geometry [12] from user inputs. However, most NeRF editing methods [24, 31, 36]




Figure 1. **Problem definition.** Given a NeRF of the original scene, and a single RGBD image of the transformed scene, we are interested in producing novel views and exporting a mesh of this transformed scene. Here we visualize the NeRF (top left) and a transformation of the scene (bottom left). We then show how the scene is re-rendered given a new camera pose in the transformed scene (top right) and its scene mesh (bottom right).

do not offer an automatic mechanism to match the transformed scene and thus require to manually define the transform (which can be non-trivial for non-rigid transformations). In our problem setting, user input is not available. Other successful works have looked at NeRF transformation through time [29, 32], where the time component is densely sampled. In contrast, we only assume a single RGBD view of the transformed scene. Although this single observation alone (without access to the original NeRF scene) can be used to directly retrieve the transformed scene via pretrained methods such as DreamGaussian [40], our experiments show that this approach struggles to recover the real geometry of the object.

Transforming a NeRF using a single RGBD introduces challenges: what is the non-rigid transformation being observed? what object parts correspond to each other? how did the unseen part (not visible in the RGBD image) deform? Inspired by mesh shape manipulation [37], we pro-

¹<https://nerfdeformer.github.io/>

pose NeRFDeformer which addresses this problem by modeling the transformation as a 3D scene flow. Concretely, the flow is a weighted linear blending of rigid transformations through 3D anchor points on the surface of the scene. This definition is more flexible than the MLP-based flow used by prior work [1, 29, 32] as we can express an approximate inverse flow. As the flow definition leverages anchor points from the original scene to the transformed scene, we design a novel robust NeRF-based correspondence matching between the NeRF scene and the RGBD observation. The method first fuses pixel correspondences from the pixel matching approach ASpanFormer [4], then applies two steps of filtering in pixel and 3D space.

We demonstrate efficacy of our method on a challenging curated dataset. This dataset was specifically designed for this problem: 113 scenes are created from 47 dynamic Objaverse assets [5]. We also propose different baselines for the problem of single-view NeRF transformation. More specifically we show that adding depth information to SINE [1] is not enough to retrieve more complicated scenes with non-rigid transformations. Our method achieves the best results for both geometric reconstruction and novel view synthesis.

Our contributions are summarized as follows: 1) We explore how a 3D scene flow can be built from 3D correspondences to transform a given NeRF to a novel scene, for which there is only single RGBD image observation. 2) We present a novel robust NeRF-based correspondence matching procedure between the original NeRF scene and the transformed observation. 3) We introduce a comprehensive new dataset for evaluating this problem setting.

2. Related Work

Neural editing and transformation. Many works have addressed neural 3D scene editing and transformation. Scene-level editing works [7–9, 16, 26, 54, 60] can change the global appearance of a scene like the global palette, style [16] or lighting [7]. This differs from object-level editing works [24, 36, 44, 52] which often learn decompositions of the scene. They can add or remove objects, or apply a rigid transformation. In general, these approaches focus on a single global rigid transformation and when present only adjust one global attribute.

Some prior works consider geometric editing. Seal-3D [46] defines the scene flow directly from a user’s 2D edits, while others [12, 31, 53, 57] use a mesh as a proxy to define local coordinates for ray bending. Importantly, the former work is only suitable for simple geometric edits like scaling or translation, while the latter works need laborious user edits in the form of 3D vertex displacements. In contrast, our method performs a non-rigid transformation given a single RGBD image and does not need laborious 3D edits.

Some conditional generative approaches [1, 13, 20, 45]

learn a distribution over NeRF parameters from a large 3D asset dataset. Editing is then formulated as mapping a given target image to a NeRF parameter [1, 20]. Such formulation restricts edits to the distribution of objects captured in the dataset, which is often not flexible enough to honor desired user requests (which we demonstrate in the experiment section). SINE [1], the closest work to ours, achieves great results on the problem of geometric editing through a single observation. Different from our formulation, SINE [1] represents flow via an MLP, which struggles to model accurate cyclic flows (forward and backward). They use FlowFormer [10] to find 2D correspondences between the transformed view and a single original view captured from the same camera pose as the transformed view. This approach limits the number of high quality correspondences. For these reasons, SINE struggles with complicated non-rigid transformations.

Notably, plenty of dynamic NeRF approaches also address the problem of deforming NeRF scenes [14, 29, 32, 50, 59]. These methods focus on deforming scenes through time, where the time component is well sampled. In contrast, we assume multiple views for one point in time and one single transformed view at a second discrete transformed time. Dynamic NeRF approaches struggle to capture the non-rigid transformation in such a setting because the amount of regularization is limited and correspondences can be hard to obtain implicitly.

Pixel correspondence matching. Optical flow methods like RAFT [43] or FlowFormer [10] predict correspondences for all pixels from an image pair. However, both are trained on image pairs with small camera movement in between which does not suit our setting. DINOv2 [28] uses self-supervised learning to learn a per-pixel embedding which can be used for correspondence matching. However, the matching is coarse and possibly less accurate to guide our 3D scene flow formulation. SuperGlue [34] matches keypoints detected from SuperPoint [6], and LoFTR [38] as well as ASpanFormer [4] match pixels using a downsampled image pair. We show that using such an approach is effective when combined with proper filtering.

Novel view synthesis from a single view. Early works on this topic conducted regression-based training on large datasets, *e.g.*, PixelNeRF [56]. Motivated by diffusion-based generative models, recent works explore how pre-trained diffusion models can aid novel view synthesis given a single view. Specifically, prior approaches exploit text-conditioned diffusion models [23, 33, 41, 51] or image-conditioned ones [18, 19, 40]. Differently, we develop an approach tailored to NeRF non-rigid transformation and obtain superior results given a pretrained NeRF.

3D scene flow representations. Prior works have studied various representations for modeling 3D scene flows. As mentioned above, many dynamic NeRF and NeRF edit-

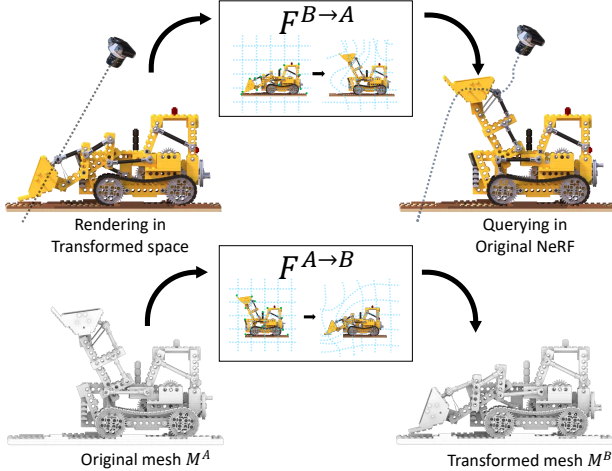


Figure 2. Overview of our method: we use two linked flows, $F^{A \rightarrow B}$ for transformed geometry reconstruction (bottom) and $F^{B \rightarrow A}$ for rendering the transformed scene (top).

ing works [1, 29, 32] apply an MLP-based flow, which works well when images are plenty. Notably, often a cyclic loss is required to connect two directions, which struggles when the transformations are complicated. Online non-rigid tracking methods [3, 11, 27] explore linear blending of anchor points as a flow design, but they do not apply their flow on NeRF-based new view synthesis. In the field of avatar modeling, many works rely on domain-specific templates, e.g., SMPL [21], to model the 3D scene flow [30, 49]. Our problem differs since we work on adapting NeRFs for general scenes and do not assume that a domain-specific object template is available.

3. NeRFDeformer

Consider an original scene A that has been transformed into a scene B , see Figure Fig. 2. Our goal is two-fold: render the transformed scene B from novel viewpoints, and extract the geometry M^B of the transformed scene. To address these goals we assume the availability of 1) a pre-trained NeRF Φ that can be used to render the original scene A from arbitrary camera poses, and 2) a single RGBD image (I^B, D^B) that captures the transformed scene B from a camera pose $C^B \in \mathbb{SE}(3)$.

At its core, our method recovers both a forward $F^{A \rightarrow B}$ and backward $F^{B \rightarrow A}$ 3D scene flow to link the two scenes:

$$p^B = F^{A \rightarrow B}(p^A) \quad (1)$$

$$p^A = F^{B \rightarrow A}(p^B), \quad (2)$$

where $p^A \in \mathbb{R}^3$ is a point in A , while $p^B \in \mathbb{R}^3$ is the corresponding point in B . As explained in the later subsections, the two transforms are defined only near the surface.




Figure 3. Forward flow of our method in the 2D case. Green dots are the anchor points v_i , the purple \times is a query point, connected to its K -nearest ($K = 3$ here) anchor points’ transformation ξ . Blue dashed lines indicate the warp of the 2D space.

Given a point p^A and direction d^A in the original scene, the NeRF can be queried for both color $c = (r, g, b)$ and density $\sigma \in \mathbb{R}_+$:

$$c, \sigma = \Phi(p^A, d^A). \quad (3)$$

To render novel views depicting the transformed scene B , we sample points (p^B) viewed along a ray in the transformed scene and apply the backward 3D scene flow $F^{B \rightarrow A}$ to obtain the corresponding points (p^A) . The direction d^A for each point is computed from the transformed difference between neighboring points along the ray, to preserve local geometry. These transformed points and directions are then fed to the original NeRF-based rendering given in Eq. (3).

Similarly, the mesh $M^A = (\mathcal{V}, \mathcal{E})$ consisting of vertices \mathcal{V} and triangle faces \mathcal{E} is obtained from the NeRF Φ of the original scene via the classic marching cubes algorithm [22]. The transformed mesh M^B is then obtained by applying the forward flow to all the original vertices:

$$M^B = (\{F^{A \rightarrow B}(v) : v \in \mathcal{V}\}, \mathcal{E}), \quad (4)$$

where we preserve topology by reusing the triangle faces \mathcal{E} .

Thus, it is apparent that the two 3D scene flows play an integral role in the process of novel view rendering, as well as in supporting recovering geometry of the transformed scene. As a result, the core of our method is aimed at recovering these scene flows. In the following we first detail the 3D scene flow is defined with local linear transformations and its trainable parameters (see Sec. 3.1). We then discuss how to optimize the trainable parameters (see Sec. 3.2) which is based on 3D corresponding points. Finally we discuss how corresponding points are extracted from the available information (see Sec. 3.3).

3.1. 3D Scene Flow as Linear Blending of Locally Rigid Transformations

In this work we define the forward flow to map any arbitrary location p^A from the original space (A) to a location p^B in the transformed space (B). The mapping is formulated

as a weighted linear blending of rigid transformations $\xi_i \in \mathbb{SE}(3)$, which are anchored at distinct 3D points. In our case the anchor points are the vertices $v_i \in \mathcal{V}$ of the triangle mesh $M^A = (\mathcal{V}, \mathcal{E})$ extracted via marching cubes from the original NeRF Φ as illustrated in Fig. 3.

Each vertex v_i has an associated 6D rigid transform ξ_i that contains a rotation R_i , a rotation origin v_i , and a translation t_i ; so that the rigid transformation and its inverse are given by

$$\xi_i(p^A) = R_i(p^A - v_i) + v_i + t_i, \quad \text{and} \quad (5)$$

$$\xi_i^{-1}(p^B) = R_i^\top(p^B - v_i - t_i) + v_i, \quad (6)$$

where each rotation matrix R_i and translation vector t_i can be thought of as parametric quantities that depend on the corresponding anchor point $v_i \in \mathcal{V}$. For more details, see the supplementary material.

The 3D flow is defined by computing the following normalized weighted sum of rigid transformations to obtain the transformed point $F^{A \rightarrow B}(p^A)$:

$$F^{A \rightarrow B}(p^A) = \sum_{k \in \mathcal{K}(p^A, \mathcal{V})} w(p^A, v_k) \xi_k(p^A), \quad (7)$$

where the summation is over K -nearest vertex neighbors, using the KNN function $\mathcal{K}(p^A, \mathcal{V})$ to return the K vertex indices that are closest to the point p^A . Each weight is defined as follows:

$$w(p^A, v) \propto \left(1 - \frac{\|v - p^A\|}{\max_{k \in \mathcal{K}(p^A, \mathcal{V})} \|v_k - p^A\|} \right). \quad (8)$$

Note that in Eq. (7) the farthest neighbor will get zero weight. The backward flow is defined as follows:

$$F^{B \rightarrow A}(p^B) = \sum_{k \in \mathcal{K}(p^B, \mathcal{V}')} w(p^B, \xi_k(v_k)) \xi_k^{-1}(p^B). \quad (9)$$

In particular, the backward flow uses transformed vertices

$$\mathcal{V}' \leftarrow \{\xi_k(v_k), \quad k = 1, \dots, |\mathcal{V}|\}. \quad (10)$$

Note from Eq. (5) that $\xi_k(v_k) = v_k + t_k$.

As mentioned in Sec. 1, our forward flow and backward flow definition has two advantages over MLP-based flows with a cyclic loss, used in prior work [1]: 1) the backward flow can be extracted from the forward flow without any training, and 2) forward and backward flows are cyclic only near the surface area where all linear transformations are similar. Thus there is no need to encourage them to be cyclic in empty space. In addition, our flow definition permits additional flexibility far from surface areas while encouraging cyclic behavior near surface areas, which is necessary for accurate geometric reconstruction and novel view synthesis.

3.2. Embedded Deformation Graph for Scene Flow Optimization

We now discuss how to find and parameterize rotation matrices R_i and translation vectors t_i for each anchor point $v_i \in \{1, \dots, |V|\}$. The optimization is inspired by embedded deformation graphs [37]. We optimize the loss

$$L_{\text{DG}} = L_{\text{ARAP}} + \alpha \cdot L_{\text{Con}} \quad (11)$$

to learn the transformation components R_i and t_i . The as-rigid-as-possible (ARAP) loss L_{ARAP} regularizes both transformation components, while the consistency term L_{Con} focuses on learning the translation terms through 3D correspondences.

The ARAP loss is applied on a decimated mesh for efficient computation. In practice when the transformation is invoked, the parametric functions R_i and t_i are computed via a weighted combination of learnable rotation matrices and translation vectors defined on the vertices of the decimated mesh. The computation is differentiable and hence end-to-end trainable. The ARAP loss regularizes the squared distances between each anchored vertex transformation applied to its neighbors and the actual transformed neighbor position. We refer the reader to the supplemental material for more details about this loss term.

The consistency loss L_{Con} constrains the translations of the vertices for which corresponding points exist. In order to ground the transformation, we first identify a set of corresponding points between scenes A and B . Let set \mathcal{I} denote the vertex indices for which correspondences exist. Thus we have the following set of corresponding points $\{(v_i^A, v_i^B) : i \in \{1, \dots, |\mathcal{I}|\}\}$. The process of selecting these points is described in the following section, with the consistency loss defined as follows:

$$L_{\text{Con}} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \|t_i + v_i^A - v_i^B\|^2. \quad (12)$$

Please note that here we do not use a rotation matrix R_i , because we are manipulating 3D vertices, and therefore no rotations are needed for transforming them. Moreover, we do not use any direct visual losses (rgb or depth) as we find reasonably dense corresponding points to suffice for learning the 3D flow.

3.3. Robust NeRF-based Correspondence Matching

We seek to produce reliable correspondences between the original NeRF scene and our transformed scene which is illustrated in a single RGBD (Fig. 4 (a)). Inspired by the work of ASpanFormer [4], we first propose to find RGB-based correspondences between the transformed RGB components and original NeRF produced renders which are filtered first in pixel space. Finally we lift the pixel correspondences to 3D and filter the false positives in 3D space.

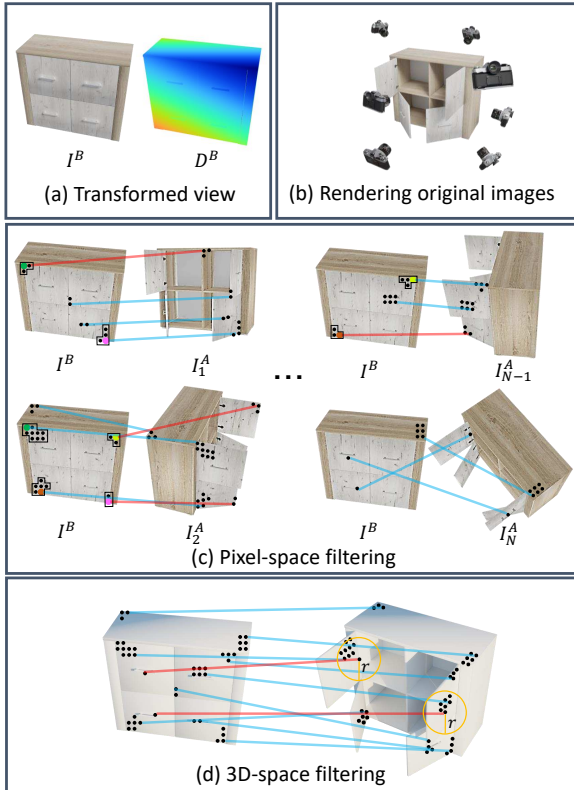


Figure 4. The transformed space image (a) is matched with the input NeRF scene first via 2D dense matching between the transformed image and original images I_1^A, \dots, I_N^A rendered from the NeRF (b). Pixel-space filtering (c) is applied where we only show selected matches (red and blue lines represent bad and good matches respectively). We show how any given pixel in I^B can be matched to multiple views (see green, yellow, and red small circles). Out of the multiple matches, we keep the one with largest continuous patch of matched pixels, *e.g.*, in I_1^A the green circle has 2 matched neighbors whereas in I_2^A there are 8. Thus we keep the latter. The points are then unprojected into 3D (d) and keep pairs that are physically close in the original space while behaving similarly in the transformed space.

2D pair correspondences and filtering. A set of images is first rendered from the provided NeRF that fully covers the hemisphere defined around the object (Fig. 4 (b)). Using ASpanFormer, dense RGB-based matching is performed between the transformed image and our RGB NeRF renders, where low confident correspondences are filtered out. To handle multiple matches between the transformed image and different renders (a given pixel might be matched to multiple locations on different images), the most confident pair of the lot is selected. This confidence is determined by the pixel neighbor density size, *e.g.*, the more the adjacent pixels have matches the more likely the matches are valid (Fig. 4 (c)). See supplemental for greater details.

3D-space filtering. Using the previous pixel correspon-

dence, their positions in 3D are lifted using the provided depth information. In order to determine which pairs are valid, points in the original scene are first clustered, and we subsequently compare how the clusters behave in the transformed scene (Fig. 4 (d)). If a cluster does not maintain its tight structure we filter the points that diverged. The intuition is as follows: point pairs that have adjacent points in the original scene should stay adjacent in the transformed scene. See supplemental for greater details.

In order to define the anchor points, \mathcal{I} , for any valid pair’s point in the original space, we find the closest vertex extracted on the mesh. This anchor point is then linked to its correspondence in 3D. Finally, these anchor point matches are used to optimize our 3D flow, as previously presented.

3.4. Implementation Details

We use $K = 20$ in the KNN employed in Eqs. (7)–(9). We only calculate the flow near the surface (surface distance $< 7e-5$) and regard other space as empty since the flow is only invertible near surface areas. We set $\alpha = 0.1$ in Eq. (11). The marching cubes resolution for M^A and mesh decimation hyperparameters are set to obtain $|\mathcal{V}| \approx 500k$ vertices. Mesh decimation is used to reduce the number of vertices to 2,000. We use Adam [15] optimizer with a learning rate of 0.001 to minimize L_{DG} for 3k iterations.

For correspondence matching, original NeRF renders are from a hemisphere which has the same distance to the object as C^B . Specifically, we sample 200 camera positions on the hemisphere, render images, and finally augment images by rotating the yaw to one of the 7 angles: $[0^\circ, -30^\circ, 30^\circ, -60^\circ, 60^\circ, -90^\circ, 90^\circ]$. More hyperparameter details are given in the appendix.

4. Experiments

Dataset. We demonstrate efficacy of baselines and our method on 113 scenes, which originate from 47 dynamic object models from the Objaverse dataset [5]. These scenes cover a wide variety of complex non-rigid transformations. For each of the 47 dynamic object models, we manually select one animation frame as the original reference and train our NeRF Φ via Instant-NGP [25] with default settings [39] for 100k iterations using 400 images with a resolution of 2880×2880 , uniformly sampled on a hemisphere above the object. Then we select one to three transformed animation frames (depending on the difficulty), different from the original animation frame(s). For each transformed time, we render one image and its corresponding depth map as the transformed view.

Baselines. We compare with generative models such as Zero123-XL [5], which finetunes a 2D diffusion model to generate new views given relative camera poses and a tar-

get image; as well as DreamGaussian [40], which is a 3D-aware Gaussian splatting based diffusion model. We also include two naive baselines: Φ and Φ finetuned. The former keeps the original NeRF without any change while the later finetunes the NeRF for an extra 2k iterations on the given transformed view using the default 2D reconstruction loss. Furthermore, we compare our method with a re-implementation (details in the supplemental) of SINE [1]. Note that for a fair comparison, for the methods that do not require depth as input like DreamGaussian, we still use the ground truth depth to solve the scale ambiguity; and SINE and our method use depth.

Metrics. For novel view synthesis, we use Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [47] and Learned Perceptual Image Patch Similarity (LPIPS) [58] as the metrics. We render 30 new views different from the training view poses and calculate the average of these metrics on 30 views, and then average across 113 scenes. For geometric reconstruction evaluation, we compute chamfer distance (CD) and Volume IoU (VmIoU). Since the process of marching cube on a collapsed NeRF scene can lead to bad reconstructions, we define a successful reconstruction when the chamfer distance is below 0.004. As such we report metrics for both all scenes (CD), for scenes that are below that threshold (CD (success)), and we also report the success rate for each method.

Results. Tab. 1 shows quantitative results for our method and baselines on our proposed dataset. See Fig. 5 for qualitative results. Tab. 1 is separated in two parts, where we first present results on visual fidelity reconstruction (left most columns) and on geometric fidelity (right most columns). The original NeRF (Φ), trained on the original scene without any further changing, performs better than other diffusion-based methods such as Zero123-XL or DreamGaussian. Diffusion-based models do not perform as well as their prior knowledge might not cover the content of the proposed scenes, *e.g.*, it has a knowledge about cats but less about doors (see Fig. 5), and it ignores the information from the prior NeRF. When fine-tuning the original NeRF scene to the transformed observation, the NeRF collapses as it does not have multiple views for constraining its behaviour. Overall our method is the best suited to both visually reconstruct the scene and extract a meaningful mesh with a success rate of 90%. We also observe that DreamGaussian [40] can capture the coarse shape of the object but lacks fine-grained texture and geometry, while SINE [1] produces inconsistent transformations as it tends to pick wrong correspondences. A real world experiment is included in Fig. 6, showing the potential of our method for handling imperfect settings (camera pose noise).

Ablations

The quantitative results of our ablation study are provided in Tab. 2. Our goal is to motivate some key design decisions in our method, especially compared with SINE. Note that the settings for the results in row 3-1 of Tab. 2 correspond to the design choices for our implementation of SINE [1]. Further note, row 4 shows the design choices of our final method.

Correspondence matching. We first analyze the effectiveness of ASpanFormer [4] as a pixel correspondence matching method and compare to FlowFormer [10] used for SINE [1]. Comparing row 1-1 and 1-2 in Tab. 2, we find ASpanFormer correspondences lead to better performance in all metrics. This is expected as ASpanFormer is trained on a dataset with large displacements while FlowFormer is trained on image pairs from adjacent frames in videos, *i.e.*, the displacements are smaller. Although the correspondences are stronger, it is still important to exploit correspondences from multiple views and filter the false positives to improve further.

Single/multiple views for correspondence. In Sec. 3.3 we presented a method that leverages NeRF to render multiple views, and as such here we evaluate the impact of object coverage (single or multiple views). Our method (row 4 in Tab. 2) is compared with the baseline in row 1-2, which only uses correspondences between the transformed image and a single original image whereas both are rendered from the same camera pose. Using multiple original images improves all metrics significantly. Fig. 7 (c,d,e) visualizes the correspondences obtained for a specific scene when using a single original image and when using multiple original images. Trivially multiple images outperforms using a single image.

Correspondence Filtering. We also compare using correspondence from multiple images obtained via ASpanFormer and only filter based on method confidence scores (row 2 in Tab. 2). Results indicate that filtering of correspondences is a non-trivial problem and our pixel-level filtering with specially designed scores and our 3D filtering are adequate for our problem (Fig. 7 (c,f)).

Scene Flow. Our scene flow is compared with the MLP cyclic flow for new view synthesis used in SINE [1]. Comparing Tab. 2 row 1-1 with row 3-1, or Tab. 2 row 4-1 with row 3-2, we observe that the MLP design hampers the performance in all metrics. We believe that the cyclic constraint is too strong and limits the expressiveness of the MLP. Replacing our flow representation method (row 4) with a MLP (row 3-2) leads to a decrease in performance. Please note that we only run experiments for visual metrics as we observed that using an MLP in spaces that do not have coverage extremely degrades the quality of the output. See Fig. 8 for qualitative results.

Depth quality. We test injecting noise to the transformed view depth via SimKinect [2]. We explore 3 different levels



Figure 5. Qualitative results comparing our method to prior work. We first show in the left-most columns the original scene and the transformed view. The other columns show different renderings of the transformed scene: ground truth in blue, DreamGaussian [40] in green, SINE [1] in yellow, and our method in red (lexicographic order within each 2×2 block).

Table 1. Main results, comparing the proposed method to baselines. CD is multiplied by 1000 for better readability.

Methods	New view synthesis			Geometric reconstruction			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	CD \downarrow	CD (success) \downarrow	succ rate \uparrow	VmIoU \uparrow
Zero123-XL [5]	14.1 \pm 3.9	0.799 \pm 0.071	0.265 \pm 0.076	/	/	/	/
DreamGaussian [40]	19.8 \pm 4.2	0.868 \pm 0.057	0.149 \pm 0.067	7.36 \pm 5.1	2.46 \pm 0.84	0.336	0.306 \pm 0.18
NeRF Φ	21.3 \pm 3.6	0.876 \pm 0.059	0.125 \pm 0.061	13.2 \pm 16	1.72 \pm 0.95	0.372	0.315 \pm 0.23
NeRF finetuned	21.6 \pm 3.5	0.826 \pm 0.096	0.198 \pm 0.100	228 \pm 270	1.85 \pm 1.10	0.195	0.312 \pm 0.25
SINE [1]*	22.1 \pm 3.8	0.883 \pm 0.052	0.115 \pm 0.053	6.40 \pm 13	1.85 \pm 1.10	0.637	0.515 \pm 0.25
Ours	25.9\pm4.2	0.924\pm0.034	0.061\pm0.040	1.46\pm2.9	0.62\pm0.79	0.903	0.666\pm0.20

Table 2. Ablation results demonstrating the efficacy of our design choices. (*FF* for FlowFormer [10], *ASpF* for ASpanFormer [4])

#	2D matching	Original images	Our flow	Filtering	New view synthesis			Geometric reconstruction		
					PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	CD \downarrow	succ rate \uparrow	VmIoU \uparrow
1-1	FF	single	\checkmark	None	22.6 \pm 4.4	0.895 \pm 0.058	0.106 \pm 0.046	6.40 \pm 13	0.637	0.515 \pm 0.25
1-2	ASpF	single	\checkmark	None	23.9 \pm 4.2	0.910 \pm 0.041	0.083 \pm 0.046	3.54 \pm 6.4	0.726	0.545 \pm 0.22
2	ASpF	multiple	\checkmark	2D naive	22.1 \pm 4.0	0.893 \pm 0.043	0.113 \pm 0.050	7.80 \pm 8.4	0.434	0.296 \pm 0.18
3-1	FF	single	MLP	None	22.1 \pm 3.8	0.883 \pm 0.052	0.115 \pm 0.053	/	/	/
3-2	ASpF	multiple	MLP	2D + 3D	22.5 \pm 3.5	0.890 \pm 0.041	0.107 \pm 0.044	/	/	/
4	ASpF	multiple	\checkmark	2D + 3D	25.9\pm4.2	0.924\pm0.034	0.061\pm0.040	1.46\pm2.9	0.903	0.666\pm0.20



Figure 6. Real world results. Left: the original scene of a half-opened box, where the first image is a training view (take from 364 images), second and third images are NeRF renders, and the last image is the mesh reconstruction. Right: the transformed scene of the same box fully opened, where the first image is the unique training view, second and third image are NeRF renders, and the last image the mesh reconstruction.

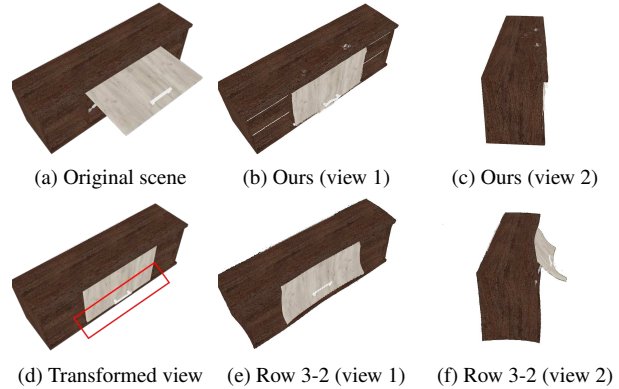


Figure 8. Rendering of different flow methods. Please refer to Tab. 2 for method references. The red rectangle in (d) highlights the region that the MLP cyclic flow struggles to model as $F^{B \rightarrow A}$ changes drastically in space. MLP flow (e,f) vs. ours (b,c).

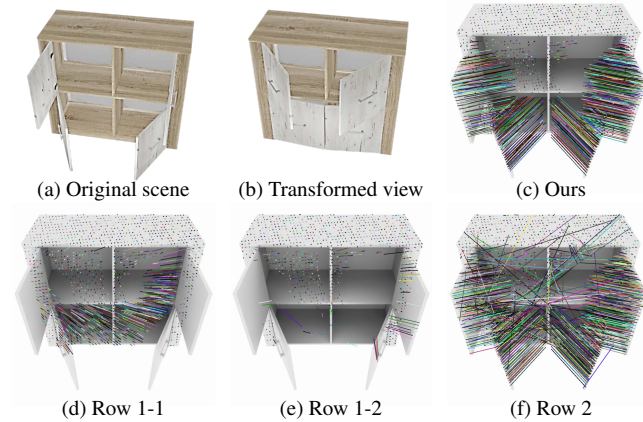


Figure 7. Correspondences. (a,b) depict the original and transformed scene. (c) to (f) are results of our method and three other ablations. Row numbers are from Tab. 2.

of noise (*s.d.* 0.3, 0.5, 1.0), where the chamfer distances are from 2.00 to 2.13 and PSNR from 23.6 to 24.4 for all of them. Depth noise does degrade the output, although our method still outperforms all baselines.

5. Conclusion

NeRFDeformer successfully transforms a NeRF scene using only a single RGBD observation of the transformed scene. The method uses local linear transformations on the surface to map the original configuration to the transformed one. In order to learn these linear transformations we introduce a new method to find dense correspondences between a NeRF scene and a single RGBD observation.

Future work should include exploring relaxing the need of depth input, such as through leveraging prior knowledge about shape or scene compositions. We are also interested in grounding diffusion models through scene flow to help determine where generation should be focused on.

Acknowledgments. Work supported in part by NSF grants 2008387, 2045586, 2106825, MRI 1725729, and NIFA award 2020-67021-32799.

References

- [1] Chong Bao, Yinda Zhang, Bangbang Yang, Tianxing Fan, Zesong Yang, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Sine: Semantic-driven image-based nerf editing with prior-guided editing field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20919–20929, 2023. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [2] Jonathan T Barron and Jitendra Malik. Intrinsic scene properties from a single rgb-d image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 17–24, 2013. [6](#)
- [3] Aljaz Bozic, Pablo Palafox, Michael Zollhöfer, Angela Dai, Justus Thies, and Matthias Nießner. Neural non-rigid tracking. *Advances in Neural Information Processing Systems*, 33:18727–18737, 2020. [3](#)
- [4] Hongkai Chen, Zixin Luo, Lei Zhou, Yurun Tian, Mingmin Zhen, Tian Fang, David Mckinnon, Yanghai Tsing, and Long Quan. Aspanformer: Detector-free image matching with adaptive span transformer. In *European Conference on Computer Vision*, pages 20–36. Springer, 2022. [2](#), [4](#), [6](#), [8](#)
- [5] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. *arXiv preprint arXiv:2307.05663*, 2023. [2](#), [5](#), [8](#)
- [6] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabbinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. [2](#)
- [7] Michelle Guo, Alireza Fathi, Jiajun Wu, and Thomas Funkhouser. Object-centric neural scene rendering. *arXiv preprint arXiv:2012.08503*, 2020. [2](#)
- [8] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, light, and material decomposition from images using monte carlo rendering and denoising. *Advances in Neural Information Processing Systems*, 35:22856–22869, 2022.
- [9] Jinkai Hu, Chengzhong Yu, Hongli Liu, Lingqi Yan, Yiqian Wu, and Xiaogang Jin. Deep real-time volumetric rendering using multi-feature fusion. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–10, 2023. [2](#)
- [10] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer: A transformer architecture for optical flow. In *European Conference on Computer Vision*, pages 668–685. Springer, 2022. [2](#), [6](#), [8](#)
- [11] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. Volumedeform: Real-time volumetric non-rigid reconstruction. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 362–379. Springer, 2016. [3](#)
- [12] Clément Jambon, Bernhard Kerbl, Georgios Kopanas, Stavros Diolatzis, George Drettakis, and Thomas Leimkühler. Nerfshop: Interactive editing of neural radiance fields. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(1), 2023. [1](#), [2](#)
- [13] Wombong Jang and Lourdes Agapito. Codenerf: Disentangled neural radiance fields for object categories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12949–12958, 2021. [2](#)
- [14] Kacper Kania, Kwang Moo Yi, Marek Kowalski, Tomasz Trzcziński, and Andrea Tagliasacchi. Conerf: Controllable neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18623–18632, 2022. [2](#)
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [16] Zhengfei Kuang, Fujun Luan, Sai Bi, Zhixin Shu, Gordon Wetzstein, and Kalyan Sunkavalli. Palettenerf: Palette-based appearance editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20691–20700, 2023. [2](#)
- [17] Yunzhi Lin, Thomas Müller, Jonathan Tremblay, Bowen Wen, Stephen Tyree, Alex Evans, Patricio A Vela, and Stan Birchfield. Parallel inversion of neural radiance fields for robust pose estimation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9377–9384. IEEE, 2023. [1](#)
- [18] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any Single Image to 3D Mesh in 45 Seconds without Per-Shape Optimization. *NeurIPS*, 2023. [2](#)
- [19] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9298–9309, 2023. [2](#)
- [20] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5773–5783, 2021. [2](#)
- [21] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM TOG*. [3](#)
- [22] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. [3](#)
- [23] Luke Melas-Kyriazi, C. Rupprecht, Iro Laina, and Andrea Vedaldi. RealFusion: 360° Reconstruction of Any Object from a Single Image. *CVPR*, 2023. [2](#)
- [24] Ashkan Mirzaei, Yash Kant, Jonathan Kelly, and Igor Gilitschenski. Laterf: Label and text driven object radiance fields. In *European Conference on Computer Vision*, pages 20–36. Springer, 2022. [1](#), [2](#)
- [25] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. [5](#)
- [26] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting

- from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8280–8290, 2022. 2
- [27] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 343–352, 2015. 3
- [28] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2
- [29] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 1, 2, 3
- [30] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural Body: Implicit Neural Representations with Structured Latent Codes for Novel View Synthesis of Dynamic Humans. In *CVPR*, 2021. 3
- [31] Yicong Peng, Yichao Yan, Shengqi Liu, Yuhao Cheng, Shanyan Guan, Bowen Pan, Guangtao Zhai, and Xiaokang Yang. Cagenerf: Cage-based neural radiance field for generalized 3d deformation and animation. *Advances in Neural Information Processing Systems*, 35:31402–31415, 2022. 1, 2
- [32] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 1, 2, 3
- [33] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, S. Tulyakov, and Bernard Ghanem. Magic123: One Image to High-Quality 3D Object Generation Using Both 2D and 3D Diffusion Priors. *ArXiv*, 2023. 2
- [34] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020. 2
- [35] William Shen, Ge Yang, Alan Yu, Jansen Wong, Leslie Pack Kaelbling, and Phillip Isola. Distilled feature fields enable few-shot language-guided manipulation. In *7th Annual Conference on Robot Learning*, 2023. 1
- [36] Karl Stelzner, Kristian Kersting, and Adam R Kosiorek. Decomposing 3d scenes into objects via unsupervised volume segmentation. *arXiv preprint arXiv:2104.01148*, 2021. 1, 2
- [37] Robert W Sumner, Johannes Schmid, and Mark Pauly. Embedded deformation for shape manipulation. In *ACM signature 2007 papers*, pages 80–es. 2007. 1, 4
- [38] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. 2
- [39] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, et al. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–12, 2023. 5
- [40] Jiayang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 1, 2, 6, 7, 8
- [41] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-It-3D: High-Fidelity 3D Creation from A Single Image with Diffusion Prior. *ArXiv*, 2023. 2
- [42] Zhenggang Tang, Balakumar Sundaralingam, Jonathan Tremblay, Bowen Wen, Ye Yuan, Stephen Tyree, Charles Loop, Alexander Schwing, and Stan Birchfield. Rgb-only reconstruction of tabletop scenes for collision-free manipulation control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1778–1785. IEEE, 2023. 1
- [43] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020. 2
- [44] Bing Wang, Lu Chen, and Bo Yang. Dm-nerf: 3d scene geometry decomposition and manipulation from 2d images. *arXiv preprint arXiv:2208.07227*, 2022. 2
- [45] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3835–3844, 2022. 2
- [46] Xiangyu Wang, Jingsen Zhu, Qi Ye, Yuchi Huo, Yunlong Ran, Zhihua Zhong, and Jiming Chen. Seal-3d: Interactive pixel-level editing for neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17683–17693, 2023. 2
- [47] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [48] Bowen Wen, Jonathan Tremblay, Valts Blukis, Stephen Tyree, Thomas Müller, Alex Evans, Dieter Fox, Jan Kautz, and Stan Birchfield. Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 606–617, 2023. 1
- [49] Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. HumanNeRF: Free-viewpoint Rendering of Moving People from Monocular Video. *CVPR*, 2022. 3
- [50] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Hu-

- mannerf: Free-viewpoint rendering of moving people from monocular video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, pages 16210–16220, 2022. 2
- [51] DeJia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. NeuraLift-360: Lifting an in-the-Wild 2D Photo to A 3D Object with 360° Views. *CVPR*, 2022. 2
- [52] 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
- [53] Bangbang Yang, Chong Bao, Junyi Zeng, Hujun Bao, Yinda Zhang, Zhaopeng Cui, and Guofeng Zhang. Neumesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In *European Conference on Computer Vision*, pages 597–614. Springer, 2022. 2
- [54] Weicai Ye, Shuo Chen, Chong Bao, Hujun Bao, Marc Pollefeys, Zhaopeng Cui, and Guofeng Zhang. Intrinsicnerf: Learning intrinsic neural radiance fields for editable novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 339–351, 2023. 2
- [55] Lin Yen-Chen, Pete Florence, Jonathan T. Barron, Tsung-Yi Lin, Alberto Rodriguez, and Phillip Isola. NeRF-Supervision: Learning dense object descriptors from neural radiance fields. In *IEEE Conference on Robotics and Automation (ICRA)*, 2022. 1
- [56] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. 2
- [57] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18353–18364, 2022. 2
- [58] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [59] Xiaoming Zhao, Alex Colburn, Fangchang Ma, Miguel Ángel Bautista, Joshua M. Susskind, and Alexander G. Schwing. Pseudo-Generalized Dynamic View Synthesis from a Video. In *ICLR*, 2024. 2
- [60] Jingsen Zhu, Yuchi Huo, Qi Ye, Fujun Luan, Jifan Li, Dianbing Xi, Lisha Wang, Rui Tang, Wei Hua, Hujun Bao, et al. I2-sdf: Intrinsic indoor scene reconstruction and editing via raytracing in neural sdf. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12489–12498, 2023. 2