

Neural Radiance Flow for 4D View Synthesis and Video Processing

Yilun Du
MIT CSAIL

Yinan Zhang
Stanford University

Hong-Xing Yu
Stanford University

Joshua B. Tenenbaum
MIT CSAIL, BCS, CBMM

Jiajun Wu
Stanford University

Abstract

We present a method, *Neural Radiance Flow (NeRFlow)*, to learn a 4D spatial-temporal representation of a dynamic scene from a set of RGB images. Key to our approach is the use of a neural implicit representation that learns to capture the 3D occupancy, radiance, and dynamics of the scene. By enforcing consistency across different modalities, our representation enables multi-view rendering in diverse dynamic scenes, including water pouring, robotic interaction, and real images, outperforming state-of-the-art methods for spatial-temporal view synthesis. Our approach works even when being provided only a single monocular real video. We further demonstrate that the learned representation can serve as an implicit scene prior, enabling video processing tasks such as image super-resolution and de-noising without any additional supervision. Code is available at <https://yilundu.github.io/nerflow/>.

1. Introduction

We live in a rich and dynamic world, consisting of scenes that rapidly change their appearance across both time and view angle. To accurately model the world around us, we need a scene representation that captures underlying lighting, physics, and 3D structure of the scene. Such representations have diverse applications: they can enable interactive exploration in both space and time in virtual reality, the capture of realistic motions for game design, and robot perception and navigation in the environment around them.

Traditional approaches, such as those used in state-of-the-art motion capture systems, typically are specialized to specific phenomena [1, 19] and fail to handle complex occlusions and fine details of motion. A core difficulty is that high resolution coverage of information requires a prohibitive amount of memory. Recent work has addressed this by using a neural network as a parametrization for scene de-

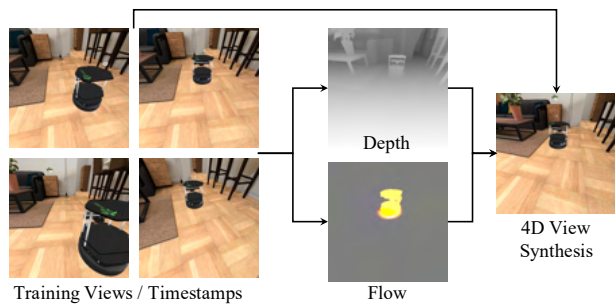


Figure 1: Given a set of training images captured from different views and timestamps, NeRFlow learns a spatial-temporal representation that captures the underlying 3D structure and dynamics and, in turn, enables 4D view synthesis.

tails [42, 64, 43]. However, these scene representation often require a static scene and a large number of images captured from many cameras, which are not generally available in real-world scenarios.

In this work, we aim to learn a dynamic scene representation which allows photorealistic novel view synthesis in complex dynamics, observed by only a limited number of (as few as one) cameras with known camera parameters. The key challenge is that the observations at each moment are sparse, restricting prior approaches [42, 64] from fitting a complex scene. To address this problem, we present a novel approach, Neural Radiance Flow (NeRFlow), that can effectively aggregate partial observations across time to learn a coherent spatio-temporal scene representation. We achieve this by formulating a radiance flow field, which encourages temporal consistency of appearance, density, and motion.

The radiance flow field is represented by two continuous implicit neural functions: a 6D (spatial position x, y, z , timestamp t and viewing direction θ, ϕ) radiance function for appearance and density, and a 4D (spatio-temporal position x, y, z, t) flow function for scene dynamics. Our representation enables joint learning of both modules, which is critical given only sparse observations at each moment. Specifically, the flow field provides temporal correspondences for spatial

locations, enabling the appearance and density information captured at different moments to propagate across time. On the other hand, the radiance function describes the scene geometry that informs the flow module about how objects are moving. Our model is fully differentiable, and thus can be trained directly using gradient backpropagation. By learning 3D structure and dynamics, our model can accomplish 4D view synthesis (Figure 1).

To evaluate our approach, we consider several challenging setups: a pouring scene which reflects fluid dynamics, an indoor scene in which a robot walks from near to far to exhibit long-range motion with great occlusion, multiple complex real scenes with transparent objects, as well as monocular videos capturing human motions. Our approach yields high-quality 4D view synthesis and outperforms a recent state-of-the-art method [43]. In addition, we show that our method can serve as a type of dynamic scene prior, which allows video denoising and super-resolution without any additional supervision, outperforming both classical and state-of-the-art internal learning methods.

In summary, our work has three contributions. First, we present a novel method, Neural Radiance Flow (NeRFlow), for learning implicit spatial-temporal scene representation. It enables novel view synthesis across both space and time. Second, we show that our approach can be effective with very limited observations down to only one camera. We achieve this by introducing a set of temporal consistency constraints over scene appearance, density, and motion. Finally, show that our approach can serve as an implicit scene prior, outperforming classical and internal learning methods in super-resolution and image de-noising.

2. Related Works

Neural scene representations. Recently, neural continuous implicit fields [9, 35, 60, 17, 79, 51, 49, 41, 64] have been developed to address the discretization issues and limited resolution of classical 3D representations such as voxel grids [5, 40, 11, 78, 58, 75, 74], point clouds [55, 13, 56, 14] and meshes [18, 71, 28, 26, 34]. Park *et al.* [51] proposed a neural signed distance function to represent scene geometry. Mescheder *et al.* [41] developed neural occupancy fields for scene reconstruction. However, they require groundtruth 3D supervision that can be difficult to obtain.

In order to learn neural scene representations directly from images, differentiable rendering [48, 23, 37, 64] is incorporated to bridge 2D observations and underlying 3D scenes. Sitzmann *et al.* [64] represented scenes with continuous feature fields and propose a neural rendering layer to allow optimization with only posed images. Niemeyer *et al.* [48] used implicit differentiation to bridge 2D images and 3D texture fields. In a recent seminal work, Mildenhall *et al.* [42] introduced a Neural Radiance Field (NeRF) that can be learned using volumetric rendering with only calibrated

images. However, these works only consider static scenes.

In contrast, we aim to learn spatial-temporal dynamic scene representations with limited observations. Although it is plausible to extend existing techniques to 4D by assuming a large number of available views at each timestep, we focus on a more realistic setting in capturing dynamic events, where only a few moving cameras are available. Our setup has significance in real-world dynamic event capturing; it also poses a great challenge in aggregating sparse, partial observations across time.

4D reconstruction. Most existing works on spatial-temporal 4D reconstruction for general scenes require sufficient observations at each moment [47, 45, 44, 30, 68, 50]. However, these methods need full observations at each timestep, and they do not recover scene appearances. Another line of work focuses on specific categories [4, 12, 21, 27, 67, 82] such as human body and faces with template models, allowing fewer observations as input. Using template models with deformations allows domain knowledge to be easily added and guarantees temporal coherence. Therefore this paradigm is widely adopted for particular shape domains [24, 4, 12, 21, 27, 67, 82] such as human face [4], body [27], and hand [59]. However, these methods depend largely on the quality of template models and it can be costly to obtain high-quality template models beyond the popular shape domains. Unlike these methods, our NeRFlow does not make domain-specific assumptions and is able to learn from limited observations.

Novel view synthesis. Although synthesizing novel views in space [42, 63, 46, 25, 20, 7, 16, 83, 73] or time (*i.e.*, video frame interpolation) [22, 39, 3, 65] is widely studied, respectively, spatial-temporal synthesis for dynamic scenes is relatively less explored [84, 32]. Recent works have extended deep learning-based novel view synthesis methods into the temporal domain, by learning a temporal warping function [36, 43] or synthesizing novel views frame-by-frame [2]. Lombardi *et al.* [36] modeled a scene by a neural feature volume and synthesize novel views at a given moment by sampling the volume with a temporally-specific warping function. Bemana *et al.* [43] targeted at view interpolation across space and time by learning a smart warping function. However, warping-based methods are restricted by input resolution. Our work is different from them in that we learn a continuous implicit representation that can theoretically scale to arbitrary resolution.

Deep networks as prior. Deep networks have been shown to manifest prior tendency for fitting natural images [69, 70, 33] and temporally-consistent videos [10], even without training on large-scale datasets. This property is referred to as an implicit image/video prior. Similarly, our method can learn neural dynamic scene representations from very sparse observations. To explain such sample efficiency, we posit that our learning method per se may serve as a ‘dynamic

implicit scene prior’. We validate it by fitting noisy and low-resolution observations, while showing good denoising and super-resolution results. Although our finding shares similar ideas with Ulyanov *et al.* [69], the prior tendency comes from our 3D rendering architecture.

Concurrent Work. Concurrent to our work, several related works [52, 77, 31, 54, 66] also investigate integrating temporal information for sparse time-step novel view synthesis. Separate from other works, we learn a single consistent continuous spatial-temporal radiance field that is constrained to generate consistent 4D view synthesis across both space and time. This enables direct rendering across both viewpoints and timestamps directly through the radiance field. This is not possible for other approaches which learn a discrete, timestamp-dependent deformation field [52, 54, 66]. Similar to our approach, [77, 31] also learn a continuous spatial-temporal radiance fields, but while our approach enforces consistency across continuous time using a neural ODE [8], they enforce consistency only at observed timestamps. Thus, while our approach can render intermediate timestamps, [77] note that interpolated renderings using their spatial-temporal radiance field are not good enough.

In addition, we further show that our approach can be applied to video processing tasks. We show that our learned radiance fields can take as input low resolution or noisy images, and can then be rendered to generate high resolution or non-noisy images.

3. Neural Radiance Flow (NeRFlow)

Our goal is to learn an implicit neural scene representation for dynamic scenes, which enables spatial-temporal novel view synthesis, given a potentially limited set of image observations and associated poses each moment. A key challenge is to effectively aggregate partial observations from different timestamps while attaining spatio-temporal coherence.

We propose Neural Radiance Flow (NeRFlow) which learns scene appearance, density and motion jointly, while encouraging temporal consistency for these components. Specifically, NeRFlow internally represents scene appearance and density by a neural radiance field, and it represents scene dynamics by a flow field. These two fields interact to propagate appearance, density, and motion information observed at different moments, modulated by a set of consistency losses derived from basic physical intuitions.

We show an overview of NeRFlow in Figure 2. In the following, we first describe the radiance and flow field. Next we present how the two fields are learned jointly using temporal consistency losses. Finally, we outline details about training supervision from RGB images and overall implementation.

3.1. Radiance and Flow Fields

NeRFlow consists of two separate modules. The first module, the radiance field, takes a 6D input of position, time

(encoded between -1 and 1), and view direction and outputs emitted color and density, which can be used to form an image by ray marching and volume rendering [42]. The second module, the flow field, takes a 4D input of position and time and outputs its flow or dynamics. We describe each function in detail below.

Radiance field. The radiance function R_θ is a 6-dimensional function, which takes as input the 4D location $\mathbf{x} = (x, y, z, t)$ and 2D viewing direction (θ, ϕ) , and outputs an emitted color $\mathbf{c} = (r, g, b)$ and volume density σ , representing the color and transparency of the corresponding 3D point (top of Figure 2). Since density is view-independent, we predict the volume density σ independently of view direction. To better aggregate cross-view visual appearance information, we also decompose the predicted color to a view-invariant diffuse part $\mathbf{c}_{\text{diffuse}}$ and a view-dependent specular part $\mathbf{c}_{\text{specular}}$. Since specularity is typically sparsely observed, during training we add an \mathcal{L}_2 regularization loss to the magnitude of $\mathbf{c}_{\text{specular}}$.

Flow field. The flow function F_θ represents the underlying dynamics of a scene. F_θ takes as input the 4D location $\mathbf{x} = (x, y, z, t)$. It outputs a flow field

$$\mathbf{f} = (f_x, f_y, f_z) = \frac{\partial x}{\partial t}, \frac{\partial y}{\partial t}, \frac{\partial z}{\partial t},$$

representing instantaneous movement of each point in space. Through integration, this function can then be used to derive the future position of any point. In particular, given a continuous point (x_s, y_s, z_s, t_s) , the future position of the point at timestamp t_g can be obtained through integration as $(x_s, y_s, z_s) + \int_{t_s}^{t_g} \mathbf{f}(x, y, z, t) dt$.

3.2. Temporally Coherent Learning

Throughout learning, we enforce consistency in radiance and flow fields so that they interact to aggregate and propagate partially observed information across time. As shown in Figure 3, this internal learning process is modulated by a set of consistency losses regarding scene appearance, density, and motion, following basic physical intuitions. Since we consider sparse observations across times that may not fully cover view angles, we only enforce consistency of diffuse color $\mathbf{c}_{\text{diffuse}}$ (Section 3.1) but not specular color.

Appearance consistency. The diffuse reflectance of an object remains constant while it is moving around. Assuming that the incidental radiance is approximately the same at the object surface [57], the emitted diffuse color remains constant. Such color constancy assumption is the basis for many optical flow algorithms and approximately holds especially when the motion is small. With this assumption, we develop the appearance consistency loss.

In particular, given a randomly sampled 3D point \mathbf{x} at timestamp t , we minimize the \mathcal{L}_2 distance between the color of \mathbf{x} and that of a predicted correspondence \mathbf{x}_c at a future

Figure 2: NeRFlow consists of two separate modules, a radiance field (top) trained via neural rendering and a flow field (bottom) trained through 3D key-point correspondence. During testing, we use only the radiance field to synthesize novel images.

timestamp \mathbf{t}_c (randomly sampled between $[\mathbf{t} - 0.5, \mathbf{t} + 0.5]$) $\mathcal{L}_{\text{RGB}} = \|\mathbf{c}_{\text{diffuse}}(\mathbf{x}) - \mathbf{c}_{\text{diffuse}}(\mathbf{x}_c)\|$, where the flow function \mathbb{R}^{F_θ} provides point correspondences, given by $\mathbf{x}_c = \mathbf{x} + \int_{\mathbf{t}}^{\mathbf{t}_c} F_\theta(x(t))dt$. This appearance consistency can be seen as a way of enabling the propagation of color information gathered in earlier (or later) timestamp to that of the current timestamp. Such propagation of color representations is especially important in settings with limited dynamic cameras where visible frames are entirely disjoint from each other across time.

Density consistency. The solidity of an object naturally remains constant while it is moving. Thus, we also enforce that density of points across time are also consistent with respect to dynamics. Analogously to the appearance consistency, we define density consistency as $\mathcal{L}_{\text{Density}} = \|\sigma(\mathbf{x}) - \sigma(\mathbf{x}_c)\|$. We note that the density consistency can be particularly useful for particles like fluid, as the shape details in fluid flowing is easily missing without some form of density consistency. Our experiment in a pouring scene shows the benefit of this.

Motion consistency. Our motion regularization is built upon two common physical intuitions: first, empty space looks static, and second, objects move smoothly in natural scenes. For the first assumption, the static empty space should consistently have no motion. Therefore, we enforce that areas with low density must exhibit low flow. To implement this, we cast N query points along a camera ray r , and select the first K query points q_k such that transmittance of the remaining camera points is greater than 0.99. We then penalize the \mathcal{L}_2 magnitude of each queried point via $\mathcal{L}_{\text{Flow}} = \|F_\theta(q_k)\|$.

The second assumption can be interpreted as that the overall scenes exhibit relatively low acceleration. Furthermore, a moving object (such as a walking robot) typically manifests similar flow at all points on its surface and within its body. Thus, we encourage the flow function to be smooth across both space and time, by penalizing the gradient of the flow functions at all randomly sampled points \mathbf{x} via $\mathcal{L}_{\text{Acc}} = \|\nabla F_\theta(\mathbf{x})\|^2$.

Figure 3: Multiple consistencies are enforced between both radiance and flow fields during training, enabling radiance information captured at earlier timestamps to inform those of later timestamps.

With these consistency losses, NeRFlow can learn a spatio-temporally coherent scene representation from limited observations at different moments.

3.3. Learning from Visual Observation

We have introduced a temporally coherent representation of NeRFlow to model dynamic scenes. We further outline training supervision from visual observations of the scene.

Volume rendering for image supervision. Given posed images (i.e. with camera matrices), we train our model using volumetric rendering following [42]. In particular, let $\{(\mathbf{c}_r^i, \sigma_r^i)\}$, denote the color and volume density of N random samples along a camera ray r . We obtain a RGB value for a pixel through alpha composition:

$$\mathbf{c}_r = \prod_{i=1}^N T_r^i \alpha_r^i \mathbf{c}_r^i, \quad T_r^i = \prod_{j=1}^{i-1} (1 - \alpha_r^j), \quad \alpha_r^i = 1 - \exp(-\sigma_r^i \delta_r^i),$$

where δ_r^i denotes the sampling distance between adjacent points on a ray. We then train our radiance function to minimize the Mean Square Error (MSE) between predicted color and ground truth RGB color via $\mathcal{L}_{\text{Render}} = \|\mathbf{c}_r - \text{RGB}\|$.

Optical flow supervision. In addition to image supervision, we also extract optical flow correspondence using Farnback’s method [15] to serve as extra supervision for predicting better scene dynamics. We then obtain sets of 3D spatial-temporal keypoint correspondences using depth maps (ground truth for synthetic scenes and through [38] for real video) and camera poses. Given 3D keypoint correspondences between points $\mathbf{x}_s = (x_s, y_s, z_s)$ observed at timestamp t_s and $\mathbf{x}_g = (x_g, y_g, z_g)$ observed at timestamp t_g , we apply integration using an Runge-Kutta solver [8] on our flow function from point \mathbf{x}_s to obtain a candidate keypoint \mathbf{x}_g^c , where $\mathbf{x}_g^c = \mathbf{x}_s + \int_{t_s}^{t_g} F_\theta(x(t))dt$. We then train our flow function to minimize the MSE between predicted and ground truth correspondences via $\mathcal{L}_{\text{Corr}} = \|\mathbf{x}_g^c - \mathbf{x}_g\|$.

3.4. Implementation Details

When training NeRFlow, we first use only $\mathcal{L}_{\text{Render}}$ to warm up training. We then train NeRFlow with our full loss $\mathcal{L}_{\text{Render}} + \mathcal{L}_{\text{Corr}} + \alpha\mathcal{L}_{\text{RGB}} + \beta\mathcal{L}_{\text{Density}} + \mathcal{L}_{\text{Flow}} + \mathcal{L}_{\text{Acc}}$, where $\alpha, \beta = 0.001$. We utilize the positional embedding [42] on the inputs to the radiance function to enable the capture of high-frequency details. We omit the positional embedding in flow functions to encourage smooth flow prediction. Please see the appendix for additional training details.

4. Experiments

We validate the performance of NeRFlow on representing dynamic scenes of pouring [61], iGibson [76], and real images from [43, 38, 80] through multi-view rendering. We further show that our approach infers high quality depth and flow maps. Finally, we show that NeRFlow can serve as a scene prior, denoising and super-resolving videos.

4.1. 4D View Synthesis

Data. We use three datasets of dynamic scenes.

Pouring: The pouring scene contains fluid dynamics [61]. We render images at 400×400 pixels. We utilize a training set size of 1,000 images and test set of 100 images.

Gibson: The Gibson scene has a robot walking a long distance. We render images on the iGibson environment [76], using the `Rs_interactive` scene with a robot TurtleBot moving linearly on the floor. Each image is rendered at 800×800 pixels. We use a training set size of 300 images and test set size of 100 images.

Real Images: Our real image datasets consists of two sources. The first are two real dynamic scenes from [43], named Ice and Vase, where Ice contains transparent objects and Vase is a complex indoor scene. The second is monocular real world videos from [38] and [80]. To evaluate on two real dynamic scenes from [43], we split 90% of provided images as the training set and the remaining 10% as the test set, and use COLMAP [62] to obtain poses for all images.

Metrics. To measure the performance of our approach, we report novel view synthesis performance using LPIPS [81], PSNR, SSIM [72], and MSE.

Baselines. We compare with four baselines. The first is the nearest neighbor baseline from Open4D [2] (using VGG feature distance). The second is a recent state-of-the-art method, X-Fields [43], which relies on warping existing training images to synthesize novel views. The third is a concurrent work, NonRigid NeRF [66], using the author’s provided codebase. We also compare with ablations.

Results on synthetic images. On synthetic images, we present a systematic analysis in three different settings for benchmarking different methods: 1) Full View, where multi-view training images are drawn uniformly across time; 2)



Images synthesized from four new views and timestamps

Figure 4: Results on Gibson in the Full View setting.

Figure 5: Illustration of cameras in the limited views setting.

Stereo Views or Dual Views, where training images are captured by two moving cameras; 3) Sparse Timestamps, where training images are drawn from a fixed, sparse subset of all timestamps in the scene.

‘Full View’ results. In the Full View setting, for the Pouring dataset, we sample cameras poses randomly in the upper hemisphere; for the Gibson dataset, we sample cameras from a set of forward facing scenes.

Tables 1 and 2 (Full View) include quantitative results on the Pouring and Gibson datasets, respectively. NeRFlow outperforms baselines in all metrics. We find that on Pouring, where modeling fluid dynamics is difficult, NeRFlow is able to capture the fluid splatter pattern and dynamics. On Gibson, which exhibits long range motion and occlusion (Figure 4), NeRFlow is able to handle of occlusions of robot.

‘Stereo Views’ and ‘Dual Views’ results. In the Stereo Views setting, training images are captured by two nearby cameras that are rotating together around a circle over time. In the Dual Views setting, training images are captured by two diametrically opposite cameras that are rotating together around a circle over time. We illustrate both settings in Figure 5. We test image synthesis from random views on any location of the circle across any time. To accomplish this task well, a model must learn to integrate the radiance information captured across different timestamps.

We report the results in Tables 1 and 2 (Stereo Views and Dual Views). Our model again outperforms all baselines. In this setting, we find that consistency enables our approach to

Models	Full View				Stereo Views				Dual Views				Sparse Timestamps			
	LPIPS↓	PSNR↑	SSIM↑	MSE↓	LPIPS↓	PSNR↑	SSIM↑	MSE↓	LPIPS↓	PSNR↑	SSIM↑	MSE↓	LPIPS↓	PSNR↑	SSIM↑	MSE↓
Nearest Neighbor	0.1023	25.34	0.9858	0.0051	0.2085	22.89	0.9667	0.0138	0.1305	25.79	0.9789	0.0088	0.1237	24.21	0.9837	0.0061
X-Fields [43]	0.0993	28.83	0.9938	0.0019	0.1261	21.25	0.9809	0.0076	0.1190	20.92	0.9787	0.0082	0.1041	28.65	0.9933	0.0021
NonRigid NeRF [66]	0.1057	31.51	0.9968	0.0009	0.1324	23.38	0.9881	0.0053	0.1057	28.12	0.9953	0.0015	-	-	-	-
NeRFlow w/o Consist.	0.1035	36.30	0.9985	0.0004	0.1219	27.98	0.9942	0.0023	0.1021	31.80	0.9982	0.0006	0.1068	33.75	0.9980	0.0006
NeRFlow (ours)	0.0980	36.57	0.9990	0.0003	0.1170	28.29	0.9958	0.0020	0.0851	35.29	0.9991	0.0003	0.0949	35.87	0.9985	0.0004

Table 1: Comparison of our approach with others on the novel-view synthesis setting on the Pouring Dataset.

Models	Full View				Stereo Views				Dual Views				Sparse Timestamps			
	LPIPS↓	PSNR↑	SSIM↑	MSE↓	LPIPS↓	PSNR↑	SSIM↑	MSE↓	LPIPS↓	PSNR↑	SSIM↑	MSE↓	LPIPS↓	PSNR↑	SSIM↑	MSE↓
Nearest Neighbor	0.1945	17.19	0.8728	0.0219	0.3314	15.03	0.8501	0.0422	0.2425	16.86	0.8832	0.0296	0.2084	16.90	0.8698	0.0250
X-Fields [43]	0.2753	21.55	0.9410	0.0096	0.3927	17.75	0.9274	0.0193	0.2587	19.13	0.9370	0.0142	0.2839	21.29	0.9378	0.0106
NonRigid NeRF [66]	0.1495	25.19	0.9616	0.0074	0.3162	19.43	0.9401	0.0132	0.2514	20.05	0.9483	0.0102	-	-	-	-
NeRFlow w/o Consist.	0.1065	29.59	0.9846	0.0028	0.2806	22.47	0.9597	0.0070	0.2729	22.26	0.9589	0.0069	0.1130	25.05	0.9712	0.0072
NeRFlow (ours)	0.0984	30.22	0.9849	0.0029	0.2496	23.65	0.9690	0.0052	0.2198	24.84	0.9758	0.0037	0.1073	25.22	0.9717	0.0070

Table 2: Comparison of our approach with others on the novel-view synthesis setting on the Gibson dataset.

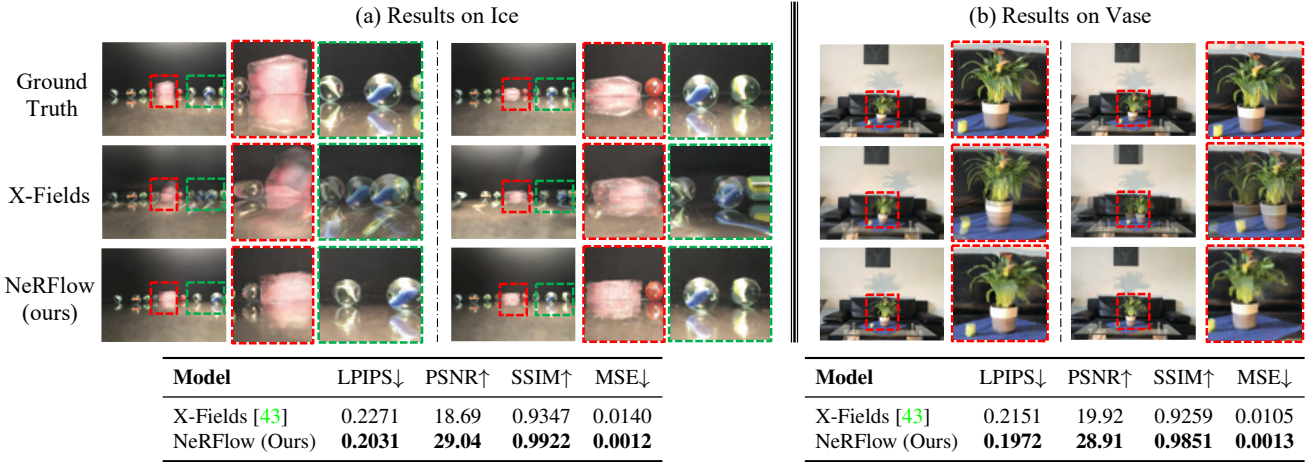


Figure 6: Novel view synthesis results on Ice (a) and Vase (b) from the sparse image datasets used by X-Fields [43].

do significantly better quantitatively, which we will elaborate later in our ablation studies in Section 4.2.

‘Sparse Timestamps’ results. We further consider the case where training images are drawn from a fixed, sparse subset of all timestamps in the scene. In particular, we train models with 1 of every 10 timestamps on Pouring, and 1 out of every 5 on Gibson. During testing, the model needs to render at arbitrary timestamps. Such a task tests the temporal interpolation capability of our model, useful for applications such as slow motion generation and frame rate up-conversion.

We report quantitative results in Tables 1 and 2 (Sparse Timestamps). NonRigid NeRF is not applicable to this setting, as it learns per timestamp latents for each timestamp in the scene. Again, NeRFlow performs well and consistency constrains radiance fields to change smoothly with respect to time, enabling smooth renderings of intermediate timestamps. Due to space constraints, qualitative results and additional analyses can be found in the supplementary material.

Results on real images.

We further evaluate our approach on real images: we

perform novel view synthesis on the image dataset in X-Fields [43], and 4D view synthesis on monocular real video datasets from [38, 80]. Figure 6 shows the novel view synthesis results on the X-Fields dataset. We find that NeRFlow captures transparency and various lighting effects in real images, while X-Fields struggles with ghosting.

We also show 4D view synthesis results in Figure 7 on monocular video datasets from [38, 80]. We visualize three different sets of results on 4D video synthesis: (1) multiple views at the same timestamp; (2) multiple timestamps from the same view; (3) randomly sampled views and timestamps. In all cases, the combination of the timestamp and the view are not in the training set. NeRFlow consistently delivers better results. Note that for the leftmost video, X-Fields appears to get incorrect poses of rendering due to a dominance of temporal warping. Additional monocular video results are in the supplementary material.

4.2. Analysis and Visualization

We next analyze NeRFlow to visualize its learned depth and flow maps, and investigate how consistency losses con-



Figure 7: Comparison of 4D synthesis (novel view and timestamp) with X-Fields on real monocular video from [80, 38].

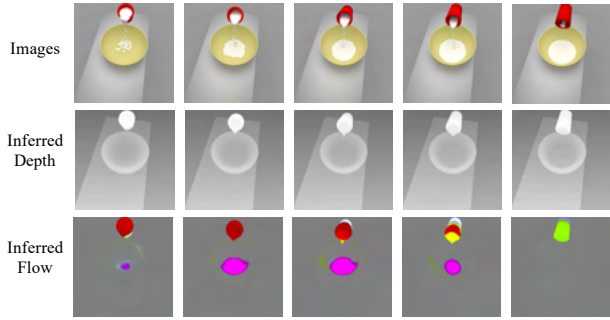


Figure 8: Visualization of estimated depth and flow (with x, y, z directions of flow represented as RGB coordinates respectively).

Models	Full View	Stereo Views	Dual Views
NeRFlow w/o Consistency	0.3747	0.4433	0.4003
NeRFlow (ours)	0.3692	0.2701	0.2675

Table 3: Evaluation of depth estimation of NeRFlow with or without physical constraints. We report MSE error with ground truth depth.

tribute to the learning of such representations and to the final results. We run our analyses on Pouring, as its simplicity leads to most interpretable results.

Visualization of depth and flow maps. In Figure 8, we visualize inferred depth and flow fields. We find that the inferred flow field captures the dynamics of pouring, including the flow of liquid as well as the movement of the cup. We quantitatively compare the depth estimation accuracy of NeRFlow with and without all consistency losses in Table 3 in terms of MSE. By enforcing geometric constancy across time, we find that our consistency loss improves depth estimates from NeRFlow, especially in limited camera settings.

Ablation study of consistency losses. Our consistency loss can be seen as a way to explicitly enforce the separation of static and dynamic components. When scene flow is predicted to be near 0, static temporal consistency of appearance/density is enforced while non-zero scene flow propa-

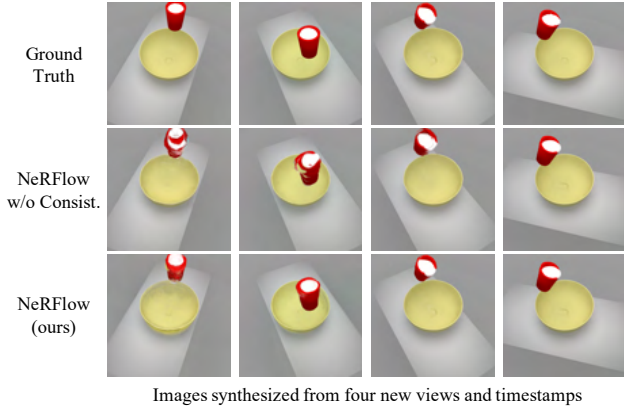
gates radiance information for dynamic modeling. We enforce zero scene flow at static locations through noisy optical flow and *motion consistency* (L_{flow} and L_{acc}).

We now ablate the effect of this separation by considering two variants: (1) reducing static separation by removing motion consistency (w/o Motion Consist.); (2) removing the dynamic modeling by only enforcing consistency on points with flow below the threshold 0.01 (w/o Dynamics Modeling). Figure 9 shows quantitative results of these models, in addition to our full and an ablated model without any consistency, on the Stereo Views setting of Pouring. Qualitatively, consistency enables more effective propagation of information across time. It makes renderings exhibit more consistent fluid placement. The supplementary material includes additional qualitative results of each variant of our ablation, which demonstrate that reducing static supervision leads to poor static structures and removing dynamic modelings leads to poor modeling of dynamic regions.

4.3. Video Processing

Given a set of images capturing a dynamic scene, NeRFlow learns to represent the underlying 3D structure and its evolution through time. This scene description can be seen as a scene prior. By utilizing volumetric rendering on our scene description, we accomplish additional video processing tasks such as video denoising and super-resolution.

Datasets. We evaluate our approach on the tasks of video denoising and image super-resolution. To test denoising, we train our model on 1,000 pouring images of the same scene with a resolution of 400×400 , rendered with a 2 ray-casts (compared with 128 rays used in Section 4.1) in Blender, and test the difference between the rendered images and the ground truth images obtained from Blender using 128 ray-casts. We also evaluate our approach on denoising a monocular real video (Ayush) from [38], where we corrupt



Models	LPIPS↓	PSNR↑	SSIM↑	MSE↓
w/o Optical Flow	0.1390	27.81	0.9932	0.0023
w/o Consist.	0.1219	27.98	0.9942	0.0023
w/o Motion Consist.	0.1372	27.85	0.9935	0.0023
w/o Dynamic Modeling	0.1317	28.09	0.9938	0.0022
NeRFlow (Full)	0.1170	28.29	0.9958	0.0020

Figure 9: Ablation study of different consistency losses on the Stereo Views setting of Pouring. Consistency regularization ensures reasonable renderings in extrapolated viewpoints. We also visualize the results of our model with or without consistency regularization.

Model	LPIPS↓	PSNR↑	SSIM↑	MSE↓
Bicubic Interpolation	0.1427	30.27	0.9961	0.0012
Blind Video Prior [10]	0.1870	30.58	0.9963	0.0009
NeRFlow (ours)	0.0903	30.67	0.9963	0.0009

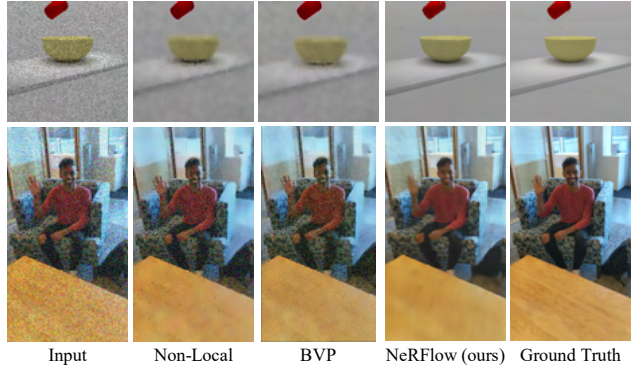
Table 4: Results of NeRFlow, Blind Video Prior [10], and bi-cubic interpolation on the task of image super-resolution.

input frames with Gaussian noise with standard deviation of 25. To test super-resolution, we train our model on 1,000 pouring images of the same scene with a resolution of 64×64 and test rendering of images of size 200×200 .

Baselines. We compare with the very recent, state-of-the-art internal learning method, Blind Video Prior [10], which uses a learned network to approximate a task mapping. During training, we supervise the Blind Video Prior on denoising and super-resolution using the outputs of the classical algorithms: Non-Local Means [6] for denoising and bi-cubic interpolation for super-resolution. We also compare with these classical algorithms directly.

Video denoising. Figure 10 shows the results on denoising. NeRFlow achieves more realistic images than the baselines and a lower reconstruction error. By accumulating radiance information over input images, our representation learns to remove most image noises. On the real monocular video, NeRFlow also obtains more realistic images than our baselines (as in LPIPS) and achieves a lower MSE.

Video super-resolution. We finally evaluate our approach on image super-resolution with our baselines in Table 4. We find that in this setting our approach again achieves more realistic images than our baselines (as determined by LPIPS),



Data	Models	LPIPS↓	PSNR↑	SSIM↑	MSE↓
Pouring	Non-Local Means [6]	0.4662	24.91	0.9263	0.0032
	Blind Video Prior [10]	0.5572	18.24	0.8891	0.0151
	NeRFlow (ours)	0.3556	28.46	0.9837	0.0014
Ayush [38]	Non-Local Means [6]	0.3051	23.49	0.9856	0.0046
	Blind Video Prior [10]	0.2707	21.67	0.9797	0.0070
	NeRFlow (ours)	0.1372	27.71	0.9949	0.0018

Figure 10: Results of NeRFlow, Blind Video Prior [10], and Non-Local Means [6] on the task of denoising.

with the Blind Video Prior achieving comparable image MSE. When rendering higher resolution images from our radiance function, representations in NeRFlow have accumulated radiance information across different input images, and are capable of rendering higher resolution details, despite being only trained on low resolution images. The supplementary material includes qualitative results.

5. Discussion

We have presented NeRFlow, a method that learns a powerful spatial-temporal representation of a dynamic scene. We have shown that NeRFlow can be used for 4D view synthesis from limited cameras (*e.g.*, monocular videos) on multiple datasets. We have also shown that NeRFlow can serve as a learned scene prior, which can be applied to video processing tasks such as video de-noising and super-resolution.

Limitations: Representing a dynamic 3D scene for view synthesis from limited image observations poses great challenges besides information aggregation. Our approach does not explicitly address the ambiguities in both 3D geometry and dynamic regions. Such ambiguity leads to difficulty in modeling complex real scenes and in preserving static backgrounds over time. We envision addressing these two challenges can greatly improve our approach, *e.g.*, explicitly separating static background and dynamic foregrounds to determine which regions should have non-zero flows, and leveraging dense depth maps to resolve geometry ambiguity.

Acknowledgements: Yilun Du is funded by an NSF graduate fellowship. This work is in part supported by ONR MURI N00014-18-1-2846 and IBM Thomas J. Watson Research Center CW3031624.

References

- [1] Bradley Atcheson, Ivo Ihrke, Wolfgang Heidrich, Art Tevs, Derek Bradley, Marcus Magnor, and Hans-Peter Seidel. Time-resolved 3d capture of non-stationary gas flows. *ACM transactions on graphics (TOG)*, 27(5):1–9, 2008. 1
- [2] Aayush Bansal, Minh Vo, Yaser Sheikh, Deva Ramanan, and Srinivasa Narasimhan. 4d visualization of dynamic events from unconstrained multi-view videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5366–5375, 2020. 2, 5
- [3] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3703–3712, 2019. 2
- [4] Amit Bermano, Thabo Beeler, Yeara Kozlov, Derek Bradley, Bernd Bickel, and Markus Gross. Detailed spatio-temporal reconstruction of eyelids. *ACM Transactions on Graphics (TOG)*, 34(4):1–11, 2015. 2
- [5] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236*, 2016. 2
- [6] Antoni Buades, Bartomeu Coll, and J-M. Morel. A non-local algorithm for image denoising. In *2005 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 8
- [7] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432, 2001. 2
- [8] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *NeurIPS*, 2018. 3, 4, 14
- [9] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. 2
- [10] Qifeng Chen, Chenyang Lei, Yazhou Xing. Blind video temporal consistency via deep video prior. <https://arxiv.org/pdf/2010.11838.pdf>, 2020. 2, 8
- [11] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *ECCV*, 2016. 2
- [12] Huseyin Coskun, Felix Achilles, Robert DiPietro, Nassir Navab, and Federico Tombari. Long short-term memory kalman filters: Recurrent neural estimators for pose regularization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5524–5532, 2017. 2
- [13] Gil Elbaz, Tamar Avraham, and Anath Fischer. 3d point cloud registration for localization using a deep neural network auto-encoder. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4631–4640, 2017. 2
- [14] Haoqiang Fan, Hao Su, and Leonidas Guibas. A point set generation network for 3d object reconstruction from a single image. In *CVPR*, 2017. 2
- [15] Gunnar Farnebäck. Two-frame motion estimation based on polynomial expansion. In *Scandinavian conference on Image analysis*, pages 363–370. Springer, 2003. 4
- [16] John Flynn, Michael Broxton, Paul Debevec, Matthew Duvall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2367–2376, 2019. 2
- [17] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7154–7164, 2019. 2
- [18] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *CVPR*, 2018. 2
- [19] Tim Hawkins, Per Einarsson, and Paul Debevec. Acquisition of time-varying participating media. *ACM Transactions on Graphics (TOG)*, 24(3):812–815, 2005. 1
- [20] 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
- [21] Yinghao Huang, Federica Bogo, Christoph Lassner, Angjoo Kanazawa, Peter V Gehler, Javier Romero, Ijaz Akhter, and Michael J Black. Towards accurate marker-less human shape and pose estimation over time. In *2017 international conference on 3D vision (3DV)*, pages 421–430. IEEE, 2017. 2
- [22] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation, 2018. 2
- [23] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1251–1261, 2020. 2
- [24] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *CVPR*, 2018. 2
- [25] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–10, 2016. 2
- [26] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. *arXiv:1803.07549*, 2018. 2
- [27] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5614–5623, 2019. 2
- [28] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *CVPR*, 2018. 2
- [29] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 14

- [30] Vincent Leroy, Jean-Sébastien Franco, and Edmond Boyer. Multi-view dynamic shape refinement using local temporal integration. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3094–3103, 2017. 2
- [31] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *arXiv preprint arXiv:2011.13084*, 2020. 3, 13
- [32] Christian Lipski, Christian Linz, Kai Berger, Anita Sellent, and Marcus Magnor. Virtual video camera: Image-based viewpoint navigation through space and time. In *Computer Graphics Forum*, 2010. 2
- [33] Jiaming Liu, Yu Sun, Xiaojian Xu, and Ulugbek S Kamilov. Image restoration using total variation regularized deep image prior. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7715–7719. IEEE, 2019. 2
- [34] Shichen Liu, Weikai Chen, Tianye Li, and Hao Li. Soft rasterizer: Differentiable rendering for unsupervised single-view mesh reconstruction. *arXiv preprint arXiv:1901.05567*, 2019. 2
- [35] Shichen Liu, Shunsuke Saito, Weikai Chen, and Hao Li. Learning to infer implicit surfaces without 3d supervision. In *Advances in Neural Information Processing Systems*, pages 8295–8306, 2019. 2
- [36] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. 2
- [37] Matthew M Loper and Michael J Black. Opendr: An approximate differentiable renderer. In *European Conference on Computer Vision*, pages 154–169. Springer, 2014. 2
- [38] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 39(4), 2020. 4, 5, 6, 7, 8
- [39] Dhruv Mahajan, Fu-Chung Huang, Wojciech Matusik, Ravi Ramamoorthi, and Peter Belhumeur. Moving gradients: a path-based method for plausible image interpolation. *ACM Transactions on Graphics (TOG)*, 28(3):1–11, 2009. 2
- [40] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *IROS*, 2015. 2
- [41] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 2
- [42] 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. *arXiv preprint arXiv:2003.08934*, 2020. 1, 2, 3, 4, 5, 14
- [43] Bemana Mojtaba, Myszkowski Karol, Seidel Hans-Peter, and Ritschel Tobias. X-fields: Implicit neural view-, light- and time-image interpolation. *ACM TOG*, 39(6), 2020. 1, 2, 5, 6, 13
- [44] Armin Mustafa, Hansung Kim, Jean-Yves Guillemaut, and Adrian Hilton. General dynamic scene reconstruction from multiple view video. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 900–908, 2015. 2
- [45] Armin Mustafa, Hansung Kim, Jean-Yves Guillemaut, and Adrian Hilton. Temporally coherent 4d reconstruction of complex dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4660–4669, 2016. 2
- [46] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7588–7597, 2019. 2
- [47] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5379–5389, 2019. 2
- [48] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. 2
- [49] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4531–4540, 2019. 2
- [50] Martin Ralf Oswald, Jan Stühmer, and Daniel Cremers. Generalized connectivity constraints for spatio-temporal 3d reconstruction. In *European Conference on Computer Vision*, pages 32–46. Springer, 2014. 2
- [51] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 2
- [52] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B. Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Deformable neural radiance fields. *arXiv preprint arXiv:2011.12948*, 2020. 3
- [53] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019. 14
- [54] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. *arXiv preprint arXiv:2011.13961*, 2020. 3, 13
- [55] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, 2017. 2

- [56] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017. 2
- [57] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001. 3
- [58] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *CVPR*, 2017. 2
- [59] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (ToG)*, 36(6):245, 2017. 2
- [60] 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*, pages 2304–2314, 2019. 2
- [61] Connor Schenck and Dieter Fox. Towards learning to perceive and reason about liquids. In *International Symposium on Experimental Robotics*, pages 488–501. Springer, 2016. 5
- [62] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [63] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019. 2
- [64] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, pages 1121–1132, 2019. 1, 2
- [65] 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 (ECCV)*, pages 155–171, 2018. 2
- [66] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video, 2020. 3, 5, 6
- [67] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *NIPS*, 2017. 2
- [68] Ali Osman Ulusoy, Octavian Biris, and Joseph L Mundy. Dynamic probabilistic volumetric models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 505–512, 2013. 2
- [69] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9446–9454, 2018. 2, 3
- [70] Dave Van Veen, Ajil Jalal, Mahdi Soltanolkotabi, Eric Price, Sriram Vishwanath, and Alexandros G Dimakis. Compressed sensing with deep image prior and learned regularization. *arXiv preprint arXiv:1806.06438*, 2018. 2
- [71] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. *arXiv:1804.01654*, 2018. 2
- [72] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. 5
- [73] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7467–7477, 2020. 2
- [74] Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. Learning a Probabilistic Latent Space of Object Shapes via 3D Generative-Adversarial Modeling. In *NIPS*, 2016. 2
- [75] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015. 2
- [76] Fei Xia, William B Shen, Chengshu Li, Priya Kasimbeg, Micael Edmond Tchammi, Alexander Toshev, Roberto Martín-Martín, and Silvio Savarese. Interactive gibbon benchmark: A benchmark for interactive navigation in cluttered environments. *IEEE Robotics and Automation Letters*, 5(2):713–720, 2020. 5
- [77] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. *arXiv preprint arXiv:2011.12950*, 2020. 3
- [78] Haozhe Xie, Hongxun Yao, Xiaoshuai Sun, Shangchen Zhou, and Shengping Zhang. Pix2vox: Context-aware 3d reconstruction from single and multi-view images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2690–2698, 2019. 2
- [79] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In *Advances in Neural Information Processing Systems*, pages 492–502, 2019. 2
- [80] Jae Shin Yoon, Kihwan Kim, Orazio Gallo, Hyun Soo Park, and Jan Kautz. Novel view synthesis of dynamic scenes with globally coherent depths from a monocular camera, 2020. 5, 6, 7
- [81] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep networks as a perceptual metric. In *CVPR*, 2018. 5
- [82] Qian Zheng, Xiaochen Fan, Minglun Gong, Andrei Sharf, Oliver Deussen, and Hui Huang. 4d reconstruction of blooming flowers. In *Computer Graphics Forum*, 2017. 2
- [83] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *ECCV*, 2016. 2
- [84] C Lawrence Zitnick, Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. High-quality video

view interpolation using a layered representation. *ACM TOG*, 23(3):600–608, 2004. [2](#)

Appendix: Neural Radiance Flow for 4D View Synthesis and Video Processing

We first present additional qualitative visualizations in Section A on both real and synthetic images. We further describe training details in Section B and provide pseudocode for the overall algorithm. We recommend reviewers to see attached videos for qualitative visualizations of our approach.

A. Additional Visualizations

Real Image View Synthesis: We illustrate novel view renderings in scenes with greater degrees of underlying dynamics from [31] and [54] in Figure 11. Compared with NonRigid NeRF, our renderings are sharper and capture dynamics more accurately.



Figure 11: Novel view renderings on dynamic scenes. Note that NonRigid NeRF collapsed to an earlier frame in the first row’s last image, and our NeRF performs better. See the video for details.

Real Depth Maps: We visualize depth images from a variety of real images in Figure 12. We find that NeRF is able to reliably infer the depth of a variety of images.

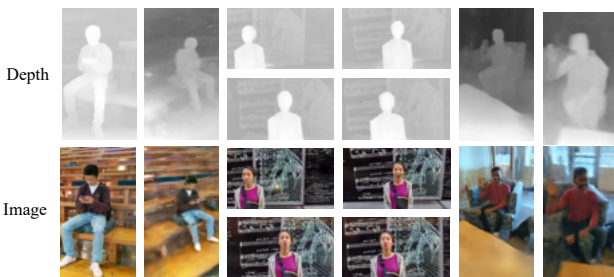


Figure 12: Illustrations of depth predicted by NeRF. NeRF is able to reliably capture the depth of dynamic scenes.

Image Interpolations: We provide visualizations of temporal interpolations of NeRF. NeRF is trained with 1 in 5 frames in the pouring scene, and rendered across all frames in a scene. With consistency, we find that we are able to consistently model drops of liquid volume throughout the duration of the pouring animation (Figure 15).

Full View Synthesis: We provide full view synthesis results on pouring in Figure 13. Compared to X-Fields [43], we

find NeRF is able to generate more coherent animations of pouring.

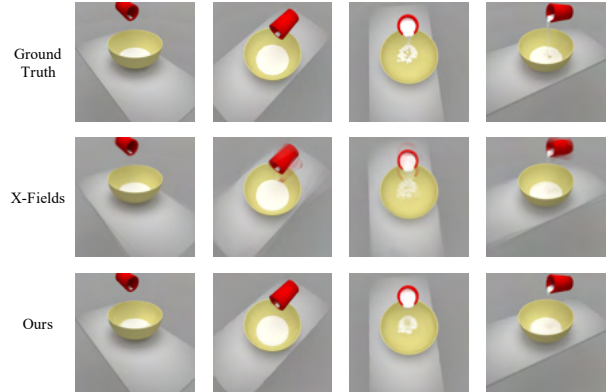


Figure 13: Results on Pouring in the Full View setting.

Ablation Visualizations: We provide visualizations of renderings of each ablation of NeRF on the stereo capture setting in pouring scene in Figure 14. With either no consistency terms or no motion consistency, the renderings have poor bowl structure. Without dynamic modeling, we observe that the underlying motion of the cup is blurred across renderings. Finally, NeRF obtains non-blurry renderings that also capture the underlying bowl structure.

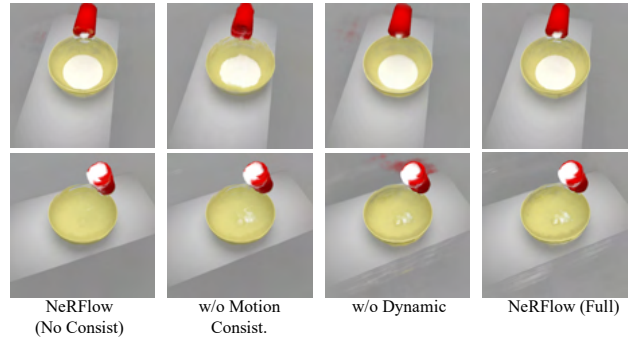


Figure 14: Qualitative comparisons of ablations. No consistency and w/o motion consistency have poor bowl structure. Without dynamic modeling, we observe that underlying motion of the cup is blurred across renderings. Finally, NeRF obtains non-blurry renderings that also capture the underlying bowl structure.

Supplemental Visualizations: We recommend readers to examine our attached supplemental video consisting of visualizations of NeRF. We first show 4D view synthesis on captured monocular videos. Next, we provide a visualization showing integration of flow across time. We then show examples of free view synthesis as well as view synthesis under stereo and dual camera configurations on Pouring and Gibson scenes. Finally, we show NeRF applied to video processing tasks of de-noising.

Figure 15: Illustration of rendering results on intermediate timestamps when models are trained with a sparse number of timestamps. Consistency makes pouring volume significantly more stable.

B. Additional Experimental Details

We train models using the PyTorch framework [53], and utilize the same base architecture as [42] for both radiance and flow functions. Models are trained for 10 hours with only $\mathcal{L}_{\text{Render}}$, and then trained for another 10 hours with all losses utilizing a single Nvidia 2080 Ti GPU. Models are trained with a learning rate of 0.001, with a exponential learning rate decay by a factor 0.1 every 40,000 steps using the Adam optimizer [29].

To integrate flow for temporal correspondences, we utilize the Neural ODE library [8], using the Runge-Kutta solver, with an $\text{RTOL} = 10^{-4}$ and $\text{ATOL} = 10^{-5}$. While these values are higher than typical values used during Neural ODE training, we found that this is critical for stable flow inference. Smaller values cause much slower training times and lead to less smooth flow fields. When training models, we penalize the predictions of $\mathbf{c}_{\text{specular}}$ using an \mathcal{L}_2 coefficient of 0.1. We apply coefficients of 0.001 for $\mathcal{L}_{\text{Corr}}$ and $\mathcal{L}_{\text{Density}}$. We provide pseudocode for training in Algorithm 1

Algorithm 1 NeRFlow Training and Sampling Algorithm

Input: Radiance function R_θ , Flow function F_θ , Spatial Correspondences $(\mathbf{x}_s, t_s), (\mathbf{x}_g, t_g)$, Camera Rays C , RGB values R , Timesteps t

▷ *Train NeRFlow Model:*

while not converged **do**

$C_i, R_i, t_i, (\mathbf{x}_s^i, t_s^i), (\mathbf{x}_g^i, t_g^i) \leftarrow C, R, t, (\mathbf{x}_s, t_s), (\mathbf{x}_g, t_g)$

$\mathcal{L}_{\text{Render}} = (\text{Render}(R_\theta, t, C_i) - R_i)^2$ # [41]

$\mathcal{L}_{\text{Flow}} = (\hat{\mathbf{x}}_g^i - \mathbf{x}_g^i)^2$ # $\hat{\mathbf{x}}_g^i$ computed from flow integration

 ▷ *Compute new spatial-temporal correspondance using F_θ :*

$(\mathbf{x}'_s, t'_s), (\mathbf{x}'_g, t'_g) \leftarrow F_\theta$

$\mathcal{L}_{\text{Density}} = \|\sigma(\mathbf{x}'_s, t'_s) - \sigma(\mathbf{x}'_g, t'_g)\|$

$\mathcal{L}_{\text{RGB}} = \|\mathbf{c}(\mathbf{x}'_s, t'_s) - \mathbf{c}(\mathbf{x}'_g, t'_g)\|$

Optimize all losses using Adam:

end while

▷ *Render From NeRFlow Model:*

$C \leftarrow V, t$ # Choose viewpoint V and timestamp t

Image = $\text{Render}(R_\theta, t, C)$ # Volumetric rendering [41]
