

FlowIBR: Leveraging Pre-Training for Efficient Neural Image-Based Rendering of Dynamic Scenes

Marcel Büsching¹

Josef Bengtson²

David Nilsson²

Mårten Björkman¹

¹KTH Royal Institute of Technology

²Chalmers University of Technology

Abstract

We introduce *FlowIBR*, a novel approach for efficient monocular novel view synthesis of dynamic scenes. Existing techniques already show impressive rendering quality but tend to focus on optimization within a single scene without leveraging prior knowledge, resulting in long optimization times per scene. *FlowIBR* circumvents this limitation by integrating a neural image-based rendering method, pre-trained on a large corpus of widely available static scenes, with a per-scene optimized scene flow field. Utilizing this flow field, we bend the camera rays to counteract the scene dynamics, thereby presenting the dynamic scene as if it were static to the rendering network. The proposed method reduces per-scene optimization time by an order of magnitude, achieving comparable rendering quality to existing methods — all on a single consumer-grade GPU.

1. Introduction

Novel view synthesis for dynamic scenes allows for rendering views of an observed scene from new viewpoints, possibly also at new points in time. Recent methods [22, 23, 32, 34] already show impressive rendering qualities. Nevertheless, they suffer from long training times and show limitations for fast-changing scenes with sparse observations [13]. We assume that these limitations are partially due to the fact that these methods are only optimized per scene without exploiting any prior knowledge. Therefore, we present *FlowIBR* which combines a per-scene learned scene flow field with a pre-trained generalizable novel view synthesis method [44] as a rendering backbone.

For static scenes there is already a multitude of generalizable novel view synthesis methods [25, 39, 42, 44, 53] which are trained on many different scenes to learn how to aggregate information from observations to synthesize novel views. Importantly, generalizable methods can therefore render novel views of previously unseen scenes *without* any scene-specific training. These methods often outperform per-scene optimized methods in terms of render-

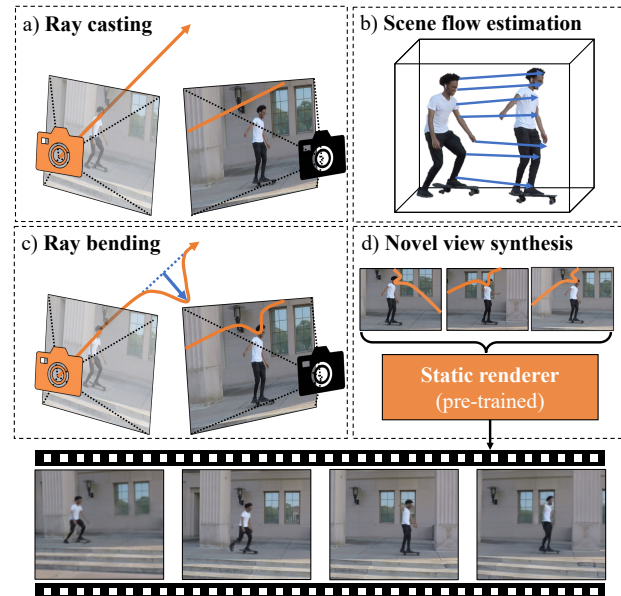


Figure 1. **Method overview** a) An image at an arbitrary position (orange camera) is synthesised based on existing observations (black camera), collected at different times. Problem: Due to the movement of the skater, the skater is not on the epipolar line of the camera ray. b) We model scene motion using per-scene learned scene flow. c) Scene flow is used to compensate the motion by bending the camera ray. d) A pre-trained neural IBR method for static scenes [44] is used for image synthesis.

ing quality, training times and demand for dense observation of the scene [25, 53, 57]. However, due to the limited availability of datasets of dynamic scenes for training, such generalized methods cannot readily be extended to dynamic scenes.

To overcome this limitation, we utilize a pre-trained rendering backbone based on Generalizable NeRF Transformer (GNT) [44], a generalizable view synthesis method which has been pre-trained on a large corpus of more easily obtainable static scenes. GNT performs neural image-based rendering (IBR) [53], synthesizing a novel view by pixel-wise projecting camera rays as epipolar lines into neighbouring

source views and then aggregating image information along them. However, in dynamic scenes, projections of independently moving scene content are likely displaced with respect to the epipolar lines. Therefore, we use a per-scene learned scene flow field to bend camera rays, so they follow the motion of the scene content over time as shown in Fig. 1. The core optimization is thus simplified by shifting focus from jointly learning the complex interplay of scene geometry, color, and temporal changes to a separate step of learning the scene dynamics.

With this approach, we are able to reduce the necessary per-scene setup time to about 1.5 hours, an improvement by an order of magnitude compared to previous methods, while getting comparable rendering quality. Additionally, this dynamics-focused optimization process allows us to create a training regime which enables optimizing the scene flow network on a single Nvidia 3080 RTX GPU. In summary, our contributions are:

- *Introducing FlowIBR*, a novel view synthesis method for dynamic scenes which combines a per-scene optimized scene flow field with a pre-trained neural image-based rendering method to decrease training time.
- *Presenting a dynamics-centred training regime* which allows for fast training of the proposed method on a single consumer-grade GPU.
- *Demonstrating the performance of FlowIBR* on the Nvidia Dynamic Scenes Dataset [56] showing competitive rendering quality despite significantly lower training times with respect to the state-of-the-art.

Code for the proposed method will be made available¹.

2. Related Work

Novel view synthesis for static scenes: Neural Radiance Fields (NeRFs) [28] and succeeding methods [1, 2, 29, 59] are able to synthesize photo-realistic images of a scene by modelling a continuous radiance and density function of an observed scene with a multi-layer perceptron (MLP). Gaussian Splatting [18] represents the scenes with a set of 3D Gaussians, which are projected on the image plane and then aggregated to a pixel color via GPU optimized alpha blending. NeRF and Gaussian Splatting based methods are optimized per scene without leveraging prior knowledge. In contrast, techniques for generalizable novel view synthesis are trained across multiple scenes, in order to allow for synthesis of novel views from sparsely observed scenes, unseen during training. Earlier methods such as pixelNeRF [57] and MVNeRF [5] have achieved this by deploying a generalizable NeRF conditioned on latent vectors extracted from the source observations. Many of the current methods [8, 17, 25, 36, 42–44, 53] combine transformers [6, 48] with multi-view geometry. As proposed by IBRNet [53], this is

realized by projecting points sampled along the camera ray for each pixel into the source observations. The information at these projected points is then extracted and subsequently aggregated into pixel color using transformers — often in combination with volumetric rendering. GNT [44] and a recent method from Du *et al.* [8] are fully transformer-based approaches, replacing the volumetric rendering, with a learned rendering.

Novel view synthesis for dynamic scenes: While early techniques required multi-view videos [21, 26, 52], current NeRF-based methods have demonstrated remarkable rendering capabilities from monocular videos. Methods such as D-NeRF [34] or Nerfies [32] extend NeRFs to the dynamic domain by learning a canonical NeRF representation at a fixed point in time [14, 32–34, 41], which is then warped by a learned motion field to other instances in time. This canonical scene representation serves as an anchor in time, accumulating the geometric and color information from source observations. A different line of work [20, 31] instantiates moving objects as separate bounded NeRFs inside a static environment NeRF, which are then rigidly moved between time-steps by learned transformations. For non-rigid dynamics, several methods [12, 22, 55] learn time-varying NeRF representations of the scene, often decomposing the scene in static and dynamic content [12, 22, 58]. To overcome the ambiguity in observing non-rigid dynamic scenes with a single camera, methods such as NSFF [22], DVS [12], NeRFlow [7] and DFNet [56] also estimate the scene flow and use it in combination with optical flow [45], depth estimates [35] and multi-view geometry as additional supervision for the per-scene learned dynamic NeRF. DVS also explores the concept of splitting the learning process — initially pre-training a static rendering framework and subsequently employing it to simplify the training of a dynamic approach. However, DVS optimizes both components only for a single scene. A recent method DynIBaR [23], which is based on NSFF, also utilizes neural IBR for novel view synthesis with techniques to address scene motion. The key distinction lies in DynIBaR being fully optimized on a single scene without utilizing pre-training, leading to long training times of up to 2 days on 8 GPUs.

Efficient view synthesis The aforementioned methods are optimized for improved visual rendering quality, but come with long per-scene training times. This problem is the object of several recent methods which introduce efficient data structures [4, 10, 11, 14, 40, 41, 50, 51] or utilize Gaussian Splatting [9, 27, 54] to reduce the computational complexity and convergence time. NeRFPlayer [41], MSTH [50] and MixVoxels [51] simplify the training process further by decomposing the scene in static and dynamic content. Although those methods show substantial improve-

¹<https://github.com/buesma/flowibr>

ments in rendering and training times, they rely on multi-view images as in the Plenoptic Video dataset [21] or quasi multi-view images [13] as in the D-NeRF dataset [34] or HyperNeRF dataset [33]. We base our method on easier to obtain monocular images, and evaluate on, according to the survey from Gao *et al.* [13], more complex Nvidia Dynamic Scenes Dataset [22, 56].

Similar to our approach, Fourier PlenOctree [52] and MonoNeRF [46] utilize generalizable novel view synthesis for dynamic scenes. However, PlenOctree also requires multi-view images that are harder to obtain than monocular videos. MonoNeRF pursues the goal of learning a fully generalizable dynamic radiance field by utilizing temporal features extracted from videos to estimate the scene flow which is then used to aggregate spatial features to pixel colors. Nevertheless, for the generalizable setting where MonoNeRF is pre-trained on one scene and then fine-tuned on the remaining scenes, only LPIPS values are reported. These are several times higher than those of the per-scene setting, for which the performance is more competitive, but at the cost of 4h of training time per scene on an A100 GPU.

3. Method

Problem formulation: Given a set of images $I_t \in \mathbb{R}^{H \times W \times 3}$ taken at discrete times $t \in \{1, \dots, T\}$ with known camera matrices $P_t \in \mathbb{R}^{3 \times 4}$, we denote the triplet (I_t, P_t, t) as an observation of the scene. We assume observations to be taken in constant time intervals Δt . The problem is now to synthesize a target view of the scene \tilde{I} from an *arbitrary* viewpoint defined by \tilde{P} and *continuous* time $\tilde{t} \in \mathbb{R}$, based on the set of available observations

$$O = \{(I_t, P_t, t)\}_{t=1}^T. \quad (1)$$

Overview: Our proposed method combines a pre-trained generalizable novel view synthesis method with a scene flow field which is learned per scene. In this section, we will review the rendering backbone used for static scenes, introduce the scene flow field model, describe the loss and regularization terms, and propose a training regime for fast training on a single GPU.

3.1. Pre-training on static scenes

For our rendering backbone we build upon Generalizable NeRF Transformer (GNT) [44], a transformer-based extension of IBRNet[53]. GNT is able to synthesize novel views of a *static* scene not seen during training, at an arbitrary target viewpoint. Novel view synthesis is performed pixel-wise in a two-stage process that involves two transformers.

View Transformer (VT): This initial stage encompasses aggregating information from source observations close to the target view. This is achieved by firstly encoding each source observation into a feature map using a U-Net [37]

image encoder $F_t = \text{U-Net}(I_t)$. Using the camera matrix of the target view, a camera ray $r_{t,m}$ is cast from the focal point e_t of the camera through the target pixel $m \in \{1, \dots, H \times W\}$ in ray direction $d_{t,m}$,

$$r_{t,m}(l) = e_t + l d_{t,m}. \quad (2)$$

Afterwards, N points $p_{t,m,n} = r_{t,m}(l_n)$ are sampled at different distances $l_n \in \mathbb{R}$ along this ray and projected onto each considered source observation. To simplify notation, we omit n and m indexes wherever possible. Consequently, these projected points are located on the respective epipolar lines, defined by the target pixel and the image planes of the source observations. Image features are extracted by interpolating the feature map at these projected points. For every individual point along the camera ray, the view transformer aggregates the associated image feature vectors through attention into a single feature vector, representing the scene content at that particular spatial location.

Ray Transformer (RT): This stage utilizes a ray transformer to conduct a learned ray-based rendering using the feature vectors created by the view transformer in the first stage. Using attention, the features along the ray are aggregated into a unified feature representation of that ray. A final multi-layer perceptron (MLP) is then applied to decode this ray feature vector, translating it into the corresponding RGB value for that particular ray.

Optimization: GNT is trained by using one of the observations as target, while employing the remaining observations as source observations. This is a technique commonly found in NeRFs [28] where a color prediction loss term

$$\mathcal{L}_{rgb} = \sum_m \|C_m - \hat{C}_m\|_2^2, \quad (3)$$

is applied to the final pixel color C_m and ground truth \hat{C}_m color of pixel m . In contrast to NeRF, GNT is trained for each training step on a different scene. This enables novel view synthesis for scenes not observed during training, without additional optimization. We select GNT as rendering backbone for our work, based on its state-of-the-art generalization capabilities and rendering quality. For more details about GNT, we refer to the original paper [44].

3.2. Scene flow field

Naively applying neural image-based rendering to dynamic scenes will yield insufficient results (see Sec. 4.3) due to the assumption that relevant neural image features can be aggregated along the epipolar lines of the target pixel on the observations. However, in dynamic scenes where each observation captures a distinct state, this faces challenges due to potential misalignment between the potentially moving scene content and the epipolar lines. Our goal is therefore

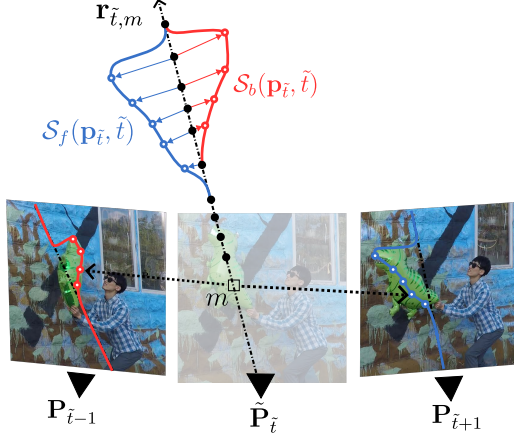


Figure 2. **Scene flow compensation** The scene flow (\mathcal{S}_f , \mathcal{S}_b) is used to adjust the ray $\mathbf{r}_{t,m}$ from the target camera $\tilde{\mathbf{P}}_t$ through the current pixel m , so that it follows the motion of the balloon at the two adjacent times $\tilde{t}+1$ and $\tilde{t}-1$. This allows the projection of the ray on the source observations to contain the pixels corresponding to m , marked by arrows.

to learn the underlying scene motion as scene flow fields [12, 22] in order to be able to realign the neural feature aggregation with the dynamic image content.

We represent per-scene learned scene flow as a forward flow field

$$\mathcal{S}_f : (\mathbf{p}_t, t) \rightarrow \mathbf{s}_f \quad (4)$$

which maps each time $t \in \mathbb{R}$ and position $\mathbf{p}_t \in \mathbb{R}^3$ to a forward scene flow vector $\mathbf{s}_f \in \mathbb{R}^3$ and a backward flow field

$$\mathcal{S}_b : (\mathbf{p}_t, t) \rightarrow \mathbf{s}_b \quad (5)$$

which maps to a backward scene flow vector $\mathbf{s}_b \in \mathbb{R}^3$. These scene flow vectors can be used to displace a point from one time to an adjacent time step with

$$\mathbf{p}_{t \rightarrow t+1} = \mathbf{p}_t + \mathcal{S}_f(\mathbf{p}_t, t) \quad (6)$$

$$\mathbf{p}_{t \rightarrow t-1} = \mathbf{p}_t + \mathcal{S}_b(\mathbf{p}_t, t) \quad (7)$$

where we generally denote $\mathbf{p}_{t \rightarrow t+1}$ as the position \mathbf{p}_t adjusted to the next time $t+1$ and $\mathbf{p}_{t \rightarrow t-1}$ to the previous time $t-1$. The fields are represented by two different heads of the same MLP

$$\text{MLP}_\theta : (\psi(\mathbf{p}), t) \rightarrow (\mathbf{s}_f, \mathbf{s}_b). \quad (8)$$

To decrease the computational complexity of estimating the scene flow, we combine a shallow MLP with the multi-resolution permutohedral lattice from PermutoSDF [38]. The number of vertices in common voxel [24] or Instant-NGP [29] based data structures grows exponentially with the input dimension, which quickly leads to a poor memory footprint and convergence time. Differently, permutohedral lattices scale linearly with the input dimension and

have been shown to work well with the 4D inputs used in our case [38]. Details of the scene flow network architecture are illustrated in the supplementary materials.

3.3. Combined method

As introduced in Fig. 2, we utilize the scene flow to adjust the 3D position of the points along the camera ray. This essentially *bends* the camera ray [34] — ensuring the ray points are projected to the parts of the image plane that contain dynamic content. Therefore, the camera rays will follow the moving scene content and thus allow static image-based rendering from being affected by scene motion, which allows for its use in dynamic scenes without any adjustments to the basic formulation.

For using more than just the two images closest to the target time \tilde{t} as source observations, it is necessary to estimate the scene flow $\mathbf{s}_{\tilde{t} \rightarrow t}$ from the target over a larger time window. This can be done with an iterative function evaluation of the scene flow network.

$$\mathbf{s}_{\tilde{t} \rightarrow t}(\mathbf{p}_{\tilde{t}}) = \begin{cases} \mathcal{S}_f(\mathbf{p}_{\tilde{t}}, \tilde{t}) + \mathbf{s}_{\tilde{t}+1 \rightarrow t}(\mathbf{p}_{\tilde{t}+1 \rightarrow t}) & , \tilde{t} < t \\ \mathcal{S}_b(\mathbf{p}_{\tilde{t}}, \tilde{t}) + \mathbf{s}_{\tilde{t}-1 \rightarrow t}(\mathbf{p}_{\tilde{t}-1 \rightarrow t}) & , \tilde{t} > t \\ 0 & , \tilde{t} = t \end{cases} \quad (9)$$

This stays computationally feasible when selecting a sufficiently small MLP (e.g. 6 to 8 layers) for the scene flow, and basing the image synthesis process on images which are taken at temporally close to the target time $\mathcal{N}(\tilde{t}) \ni t$, to limit the number of times to adjust the rays for.

Rendering at continuous target times $\tilde{t} \in \mathbb{R}$ outside the intervals Δt , in which the scene has been observed, is facilitated by estimating the scene flow at the target time, and then linearly scaling the scene flow vectors, so they displace the points not over the full interval Δt but to the times of the next observations. Following this initial step, the motion adjustment can continue as previously described. A more formal introduction of this can be found in the supplementary materials.

3.4. Losses

The problem of learning a non-rigid dynamic scene from monocular observations is highly ambiguous. For example, an object that has seemingly grown from one observation to the next can either have changed its actual size or just moved towards the camera. This ambiguity creates the need for additional supervision and regularization besides the previously introduced \mathcal{L}_{rgb} loss. For the selection of losses that are used, we take inspiration from previous scene flow-based works [12, 22, 23].

Optical flow loss: Given the absence of a ground truth for the scene flow, the optical flow between the target image and source observations serves as an effective proxy which can

be used for additional supervision. For this we use RAFT [45] to estimate the optical flow $\mathbf{o}_{\tilde{t} \rightarrow t}$ from the target image at time \tilde{t} to the used source observations at times t .

The optical flow estimate for each pixel is a 2D vector that represents the displacement to the corresponding pixel location in a second image. We use this as supervision for the 3D scene flow, by taking the scene flow-adjusted points along the target ray, and projecting them onto the image plane of the observations. Then we calculate the difference between the pixel locations of the projected points and the pixel in the target image corresponding to the camera ray. This yields one pixel displacement vector per projected ray point. We use the attention weights of the ray transformer to estimate a weighted average $\mathbf{d}_{\tilde{t} \rightarrow t}$ of these. The optical flow loss is then calculated with the L1-norm between the optical flow $\mathbf{o}_{\tilde{t} \rightarrow t}$ estimated with RAFT and the weighted average $\mathbf{d}_{\tilde{t} \rightarrow t}$ of projected scene flow,

$$\mathcal{L}_{of} = \sum_{t \in \mathcal{N}(\tilde{t})} \|\mathbf{o}_{\tilde{t} \rightarrow t} - \mathbf{d}_{\tilde{t} \rightarrow t}\|_1. \quad (10)$$

This procedure is comparable to related methods [22, 23] which weigh pixel displacements with the opacity values of the NeRF representation. Due to noise in the optical flow, we only use \mathcal{L}_{of} to initialize the scene flow to a general direction, and linearly anneal its weight during training.

Cycle consistency regularization: To ensure that the learned forward and backward scene flows are consistent to each other, we utilize *cycle consistency* regularization [12, 22, 23]. This means that the backwards scene flow of point \mathbf{p}_t should be equal to the forward scene flow of the displaced point $\mathbf{p}_{t \rightarrow t-1}$, but in the opposite direction,

$$\mathcal{L}_{cyc} = \sum_{t \in \mathcal{N}(\tilde{t})} \|\mathcal{S}_b(\mathbf{p}_t, t) + \mathcal{S}_f(\mathbf{p}_{t \rightarrow t-1}, t-1)\|_1 + \|\mathcal{S}_b(\mathbf{p}_{t \rightarrow t+1}, t+1) + \mathcal{S}_f(\mathbf{p}_t, t)\|_1. \quad (11)$$

Scene flow regularisation: We introduce three supplementary regularization terms to the scene flow, guiding it to learn anticipated properties of the inherent scene motion. For this we use the squared L2-norm between forward and backward scene flow as *temporal smoothness* regularization [12, 22, 23, 49]

$$\mathcal{L}_{temp} = \sum_{t \in \mathcal{N}(\tilde{t})} \|\mathcal{S}_f(\mathbf{p}_t, t) + \mathcal{S}_b(\mathbf{p}_t, t)\|_2^2 \quad (12)$$

which encourages the learning of a piece-wise linear solution. Additionally, we use the L1-norm on the predicted scene flow to support a generally *slow* scene flow, based on the common assumption, that most of the scene does not contain motion [12, 22, 47].

$$\mathcal{L}_{slow} = \sum_{t \in \mathcal{N}(\tilde{t})} \|\mathcal{S}_f(\mathbf{p}_t, t)\|_1 + \|\mathcal{S}_b(\mathbf{p}_t, t)\|_1 \quad (13)$$

Lastly, we encourage *spatial smoothness* so that narrow points along a ray are going to have similar scene flow [22, 23, 30]. For this, the difference between the forward and backward flows of neighbouring points $\mathbf{p}' \in \mathcal{N}(\mathbf{p})$ is weighted based on their distance

$$w(\mathbf{p}', \mathbf{p}) = \exp(-2\|\mathbf{p}' - \mathbf{p}\|_2^2). \quad (14)$$

The loss is then calculated on the L1-norm between the deviations in the flow.

$$\mathcal{L}_{spat} = \sum_{t \in \mathcal{N}(\tilde{t})} \sum_{\mathbf{p}' \in \mathcal{N}(\mathbf{p}_t)} (\|\mathcal{S}_f(\mathbf{p}_t, t) - \mathcal{S}_f(\mathbf{p}', t)\|_1 + \|\mathcal{S}_b(\mathbf{p}_t, t) - \mathcal{S}_b(\mathbf{p}', t)\|_1) \times w(\mathbf{p}', \mathbf{p}_t) \quad (15)$$

We summarize these losses to an aggregated regularization loss term

$$\mathcal{L}_{reg} = \mathcal{L}_{temp} + \mathcal{L}_{slow} + \mathcal{L}_{spat}. \quad (16)$$

Final loss: The final total loss is calculated by a weighted summation over the different loss terms

$$\mathcal{L} = \mathcal{L}_{rgb} + \alpha_{of}^k \mathcal{L}_{of} + \alpha_{cyc} \mathcal{L}_{cyc} + \alpha_{reg} \mathcal{L}_{reg}, \quad (17)$$

with α_{of}^k being linearly annealed over k steps of training.

3.5. Dynamics focused optimization

In the following, we will describe details of the process which will allow for optimization on a single GPU. For this we take advantage of the aspect that the optimization can focus on learning the scene flow, since the rendering has already been learned during pre-training. The general idea is to follow a coarse-to-fine approach [32] in learning scene flow, focusing on eliciting a coarse scene flow at the beginning, and then refining it as training time progresses.

Low-to-high number of source images: Training starts with a low number of source images to initially learn the scene flow between close time steps. Once the scene flow has started to emerge, we steadily increase the maximum number of source images at pre-defined steps to also learn more nuanced long-term dynamics. To keep a constant maximum GPU utilization, we adjust the ray batch size so that the total number of rays for all views remains constant, with a larger batch size in the beginning and then smaller batch sizes as the number of views increases.

Coarse-to-fine image resolutions: Directly training the model on the original image resolution often leads to a zero collapse of the scene flow, i.e. the scene is incorrectly predicted to be fully static. To avoid falling into a trivial local minimum created by high-frequency details, we apply

a coarse-to-fine scheme with the images originally subsampled by a factor of $f = 8$. The subsampling reduces the overall complexity of the correspondence problem, and the batches of rays are more likely to cover also small dynamic areas. During training, we incrementally increase the image resolution by reducing f in steps of 2 to gradually refine the scene flow to represent the finer details of the underlying scene dynamics.

Masked ray sampling: The regularization factors in \mathcal{L}_{reg} benefit learning of a near-zero scene flow in static regions of the scene. This allows us to focus the computational resources on the dynamic regions of the scene. We do so using binary motion masks that represent a coarse segmentation of dynamic regions in the image. Instead of sampling pixels uniformly across the image, we increase the sampling probability of pixels in regions that are potentially dynamic.

Motion masks are estimated by using an off-the-shelf instance segmentation network [15]. Of the detected instances, we retain potentially dynamic classes (e.g. person, car) with a confidence above an intentionally small threshold ($p = 0.1$) to avoid discarding dynamic content. Overly small instances are removed via morphological opening, where we use a large dilation filter mask to create a wide boundary around the remaining segments which allows for learning a better transition from static to dynamic regions. The segmentation mask also allows us to downscale the \mathcal{L}_{slow} loss at potentially dynamic pixels, encouraging a stronger scene flow. Li *et al.* [22] also utilize dynamic masks, but use them in a separate stage at the beginning of the training to sample additional rays in dynamic regions.

Fine-tuning of rendering backbone: To further increase the rendering quality we also fine-tune the rendering backbone per-scene. This is possible with minimum computational overhead, since GNT is part of the optimization loop of the scene flow network. We found this especially beneficial in combination with the optical flow loss, since the fine-tuning leads to refined attention weights of the ray transformer and therefore a more refined optical flow loss.

4. Experimental Results

In this section, we show the capabilities of our method, compare it to state-of-the-art baselines and ablations in quantitative and qualitative experiments, as well as analyse the learned scene flow.

4.1. Experimental setup

Optimization of the scene flow network and fine-tuning of GNT is done jointly with the Adam optimizer [19] on a Nvidia 3080 RTX GPU. For the GNT backbone, we train a smaller version of the original GNT architecture with 4

transformer blocks instead of 8 and otherwise follow the training regime from Varma T. *et al.* [44]. This allows for faster rendering and the reduced memory permits a higher batch-size. Further implementation and parameter details can be found in the supplementary materials.

As dataset for the evaluation we use the Nvidia Dynamic Scenes Dataset [56]. The eight scenes in the dataset are observed in a resolution of 1920×1080 by 12 static and synchronized cameras. A moving camera is simulated by selecting one of the cameras for each timestep in a circular fashion to create the train set. Then for each training step, one camera position is selected to sample a ray batch, while the source observations are selected from the set of remaining images. During evaluation, observations present in the train set are kept as the source observations, while observations unseen during training are used as target images to estimate the metrics. We utilize two versions of the dataset which differ in the temporal resolution: A *default* version with 24 timesteps, and the *long* version from Li *et al.* [23] with, depending on the scene, 90 to 200 timesteps. If not explicitly mentioned, the default version is used.

Calculating the optical flow, motion masks and subsampling for the scenes takes on average 6:45min for the default and 13:51min for the long scenes. This is similar to the set-up times of the baselines and is therefore not included in later reported training times.

For the quantitative evaluations we report Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [3] and Learned Perceptual Image Patch Similarity (LPIPS) [60]. In addition to estimating these metrics for the full image, we also estimate them for dynamic image regions, leveraging the motion masks from the dataset.

4.2. Comparative evaluation

We benchmark FlowIBR against NSFF [22], DVS [12] and HyperNeRF [33], retraining each baseline using the configurations provided by the respective authors. For DynIBaR [23], which represents the current state-of-the-art in terms of rendering quality, we present the numbers reported in the original publication. For a fair comparison in rendering time, we use images down-sampled by a factor of $f = 4$ to a resolution of 480×270 before using them for training or evaluation of the methods.

Qualitative results: Fig. 3 presents images rendered with FlowIBR and the retrained baselines. FlowIBR demonstrates the capability to synthesize novel views from previously unobserved viewpoints, achieving a quality that is competitive with the state-of-the-art. Further qualitative results are presented in the form of image sequences in the supplementary materials and as continuous renderings in the accompanying video.

Table 1. **Quantitative evaluation on Nvidia Dynamic Scenes [56]** We report metrics as average over all 8 scenes for the whole image as well as dynamic parts only. For NSFF we report results on two different training checkpoints with the number of train steps in brackets. FlowIBR (ours) has a significantly lower training time than the compared methods while obtaining competitive rendering quality.

Data	Method	Train Time			Render Time		Full Image			Dynamic Regions		
		N _{GPU}	h	GPUh	N _{GPU}	GPUsec	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
default	HyperNeRF	4	16:00	64:00	4	10.0	20.03	0.481	0.212	17.40	0.327	0.301
	DVS	4	18:30	74:00	1	20.4	25.98	0.730	0.083	22.12	0.690	0.152
	NSFF (30k)	4	2:05	8:20	1	6.2	26.57	0.604	0.103	20.84	0.627	0.178
	NSFF (1M)	4	58:15	233:00	1	6.2	28.20	0.681	0.044	21.75	0.681	0.103
	FlowIBR (Ours)	1	1:27	1:27	1	11.8	26.81	0.742	0.099	21.13	0.651	0.192
long	DynIBaR [†]	8	48:00	384:00	1	20.0	30.92	0.958	0.027	24.32	0.827	0.061
	FlowIBR (Ours)	1	1:32	1:32	1	11.3	27.32	0.793	0.083	22.06	0.702	0.132

[†] Results were taken from the original publication




Figure 3. **Qualitative evaluation on Nvidia Dynamic Scenes (default) [56]** Renderings for FlowIBR and the retrained methods (HyperNeRF, DVS, NSFF). We can clearly see that FlowIBR is able to synthesize novel views from previously unobserved viewpoints, with quality close to the ground truth (GT) image and to state-of-the-art methods.

Quantitative results: As the results in Tab. 1 indicate, FlowIBR performs competitively with respect to the baselines, but requires significantly shorter training time and only uses a single GPU. In terms of rendering quality, the results are the most comparable to those of NSFF, which requires an order of magnitude more GPU hours for training. To better understand the relative performance of the methods, we limited NSFF to a more similar training budget, causing it to underperform in comparison to FlowIBR. The moderate rendering speed of our approach can be attributed to two primary factors. First, we employ a generalizable rendering backbone that, although highly versatile, has not been specifically optimized for speed. Second, our image-based rendering approach necessitates the projection of camera rays across all source observations, unlike NeRF-based methods that often render in a single forward pass.

4.3. Ablation study

We evaluate ablations of our method, summarized in Tab. 2, to analyze the contribution of the different components. The full method outperforms all ablated versions, although some ablations show high performance in individual metrics. Most importantly, the results show that the introduction of flow compensation improves the image quality by 3.3 in PSNR for the whole image and 3.2 for the dynamic regions of the image. We found 60k training steps to be the optimal training time as displayed in Tab. 3. Shorter and longer training times lead to consistently worse metrics, especially in the dynamic parts of the image. Replacing the permutohedral encoding with an vanilla ReLU MLP yields marginally better qualitative results, at the cost of an increased rendering time of 15.1 s/img – more details on this

Table 2. **Method ablation** Ablations are obtained by selectively omitting one component at a time: (1) GNT without the scene flow, (2-3) const. subsampling f instead of the coarse-to-fine subsampling, (4) optical-flow loss \mathcal{L}_{of} , (5) cycle loss \mathcal{L}_{cyc} , (6) scene flow reg. losses $\mathcal{L}_{reg} = \mathcal{L}_{slow} + \mathcal{L}_{spat} + \mathcal{L}_{temp}$, (7) the usage of dynamic-static masks for ray sampling, and (8) the GNT fine-tuning. Metrics are averaged over the scenes *Balloon1*, *Truck* and *Playground*. Best result is **bold** and second best underlined.

Ablation	Full Image			Dynamic Regions		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
(1) default GNT	24.1	0.693	0.233	18.5	0.482	0.394
(2) $f = 4$	26.0	0.732	0.157	19.8	0.589	0.289
(3) $f = 8$	25.8	0.717	0.130	20.7	0.534	0.241
(4) \mathcal{L}_{of}	26.5	0.781	0.130	19.8	0.532	0.281
(5) \mathcal{L}_{cyc}	26.0	0.793	0.118	<u>21.2</u>	0.610	0.206
(6) \mathcal{L}_{reg}	26.3	0.795	0.124	21.0	<u>0.625</u>	<u>0.205</u>
(7) dyn. mask	26.2	0.803	0.127	19.7	0.488	0.275
(8) fine-tune	25.3	0.782	0.145	18.5	0.492	0.258
full	<u>26.4</u>	<u>0.799</u>	0.118	21.5	0.631	0.197

Table 3. **Training time ablation** Training time based parameters such as subsample frequency and learning rate decay were scaled proportionally for the ablations. Metrics are averaged over the scenes *Balloon1*, *Truck* and *Playground*.

Training steps	h	Full Image			Dynamic Regions		
		PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
30k	0:43	24.31	0.774	0.148	18.90	0.511	0.252
60k	1:29	26.42	0.799	0.118	21.53	0.631	0.195
90k	2:13	22.57	0.623	0.211	19.03	0.483	0.320

can be found in the supplementary.

4.4. Analysis of learned scene flow

To analyze the learned scene flow, we visualize the forward (s_f) and backward (s_b) scene flow for the rays of a selected reference view in Fig. 4. We project this flow onto the image plane and color code the direction and magnitude. The flow is predominantly observed at pixels corresponding to the balloon, and the color-coded directions indicate a proper counter-clockwise movement, where s_f and s_b are nearly inverse to each other. This visualization illustrates the capability of FlowIBR in capturing correct scene dynamics. Nevertheless, the scene flow does not exactly follow the contour of the balloon, which is one explanation for the occasional motion blur in rendered images.

4.5. Limitations

Since FlowIBR is initialized with optical flow, it inherits issues that complicate optical flow estimation, such as occlusions and undefined flow in large homogeneous areas. Long sequences can become challenging as a singular scene flow network, restricted by its finite capacity, must capture the entire flow. In Fig. 5 we illustrate common failure cases. These issues can typically be alleviated through scene-specific parameter tuning or alternative network initialization.

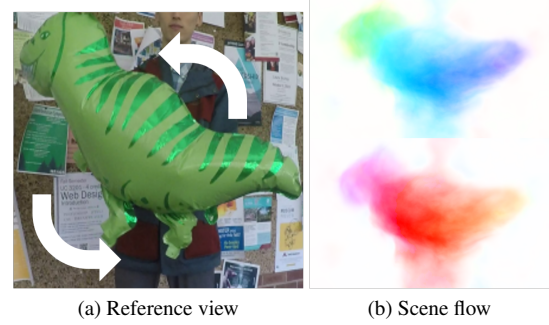


Figure 4. **Scene flow visualization** Projection of s_f (top) and s_b (bottom) onto the image plane. The arrows in the reference view show the true scene motion. Scene flow is visualized with hue indicating the direction and intensity the magnitude. Here our method correctly learned the rotating motion shown in the image.




Figure 5. **Failure cases** (a) Fast moving objects can exhibit motion blur. (b) For small objects, FlowIBR occasionally fails to learn any scene flow at all, instead learning a continuous near-zero function. (c) In some cases, images contain rendering artifacts, especially for unbounded backgrounds.

5. Discussion and Conclusion

In this paper, we presented FlowIBR, a novel view synthesis method for dynamic scenes which utilizes a pre-trained rendering method, to decrease the necessary training time. We solve the problem of a limited number of dynamic training scenes by employing a per-scene learned scene flow network, which is used to adjust the observations from different points in time, so they appear to be static to the rendering module. This allows for the utilization of a readily-available rendering backbone for static scenes in the dynamic domain – enabling shorter training times on a single consumer-grade GPU. We would like to address the moderate rendering speed of FlowIBR in future work, by creating a specialized rendering backbone which focuses even more on fast training and rendering.

Acknowledgements

This work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

- [1] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5835–5844, 2021. 2
- [2] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5460–5469, 2022. 2
- [3] Dominique Brunet, Edward R. Vrscay, and Zhou Wang. On the Mathematical Properties of the Structural Similarity Index. *IEEE Transactions on Image Processing*, 21(4):1488–1499, 2012. 6
- [4] Ang Cao and Justin Johnson. HexPlane: A Fast Representation for Dynamic Scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 130–141, 2023. 2
- [5] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. MVSNeRF: Fast Generalizable Radiance Field Reconstruction from Multi-View Stereo. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14104–14113, 2021. 2
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*, 2022. 2
- [7] Yilun Du, Yanan Zhang, Hong-Xing Yu, Joshua B. Tenenbaum, and Jiajun Wu. Neural Radiance Flow for 4D View Synthesis and Video Processing. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14304–14314, 2021. 2
- [8] Yilun Du, Cameron Smith, Ayush Tewari, and Vincent Sitzmann. Learning To Render Novel Views From Wide-Baseline Stereo Pairs. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4970–4980, 2023. 2
- [9] Bardienus P. Duisterhof, Zhao Mandi, Yunchao Yao, Jia-Wei Liu, Mike Zheng Shou, Shuran Song, and Jeffrey Ichnowski. MD-Splatting: Learning Metric Deformation from 4D Gaussians in Highly Deformable Scenes. *arXiv preprint*, arXiv:2312.00583, 2023. 2
- [10] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast Dynamic Radiance Fields with Time-Aware Neural Voxels. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 2
- [11] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-Planes: Explicit Radiance Fields in Space, Time, and Appearance. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12479–12488, 2023. 2
- [12] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic View Synthesis from Dynamic Monocular Video. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5692–5701, 2021. 2, 4, 5, 6
- [13] Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular Dynamic View Synthesis: A Reality Check. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 33768–33780, 2022. 1, 3
- [14] Xiang Guo, Jiadao Sun, Yuchao Dai, Guanying Chen, Xiaoqing Ye, Xiao Tan, Errui Ding, Yumeng Zhang, and Jingdong Wang. Forward Flow for Novel View Synthesis of Dynamic Scenes. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16022–16033, 2023. 2
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017. 6
- [16] Dan Hendrycks and Kevin Gimpel. Gaussian Error Linear Units (GELUs). *arXiv preprint*, arXiv:1606.08415, 2023. 12
- [17] Mohammad Mahdi Johari, Yann Lepoittevin, and Francois Fleuret. GeoNeRF: Generalizing NeRF with Geometry Priors. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18344–18347, 2022. 2
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023. 2
- [19] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 6, 12
- [20] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic Neural Fields: A Semantic Object-Aware Neural Scene Representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12861–12871, 2022. 2
- [21] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, and Zhaoyang Lv. Neural 3D Video Synthesis from Multi-view Video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5511–5521, 2022. 2, 3
- [22] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural Scene Flow Fields for Space-Time View Synthesis of Dynamic Scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6494–6504, 2021. 1, 2, 3, 4, 5, 6
- [23] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snavely. DynIBaR: Neural Dynamic Image-Based Rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4273–4284, 2023. 1, 2, 4, 5, 6
- [24] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural Sparse Voxel Fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 15651–15663, 2020. 4

- [25] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. Neural Rays for Occlusion-aware Image-based Rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7814–7823, 2022. 1, 2
- [26] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: learning dynamic renderable volumes from images. *ACM Transactions on Graphics*, 38(4):65:1–65:14, 2019. 2
- [27] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3D Gaussians: Tracking by Persistent Dynamic View Synthesis. In *International Conference on 3D Vision (3DV)*, 2024. 2
- [28] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *European Conference on Computer Vision (ECCV)*, pages 405–421, 2020. 2, 3, 12
- [29] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):1–15, 2022. 2, 4
- [30] Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 343–352, 2015. 5
- [31] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural Scene Graphs for Dynamic Scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2855–2864, 2021. 2
- [32] 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 *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5845–5854, 2021. 1, 2, 5
- [33] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. HyperNeRF: a higher-dimensional representation for topologically varying neural radiance fields. *ACM Transactions on Graphics*, 40(6):1–12, 2021. 3, 6
- [34] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10313–10322, 2021. 1, 2, 3, 4
- [35] Rene Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-Shot Cross-Dataset Transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3):1623–1637, 2022. 2
- [36] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common Objects in 3D: Large-Scale Learning and Evaluation of Real-life 3D Category Reconstruction. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10881–10891, 2021. 2
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015)*, pages 234–241, 2015. 3
- [38] Radu Alexandru Rosu and Sven Behnke. PermutoSDF: Fast Multi-View Reconstruction with Implicit Surfaces Using Permutohedral Lattices. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8466–8475, 2023. 4, 12
- [39] Mehdi S.M. Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lucic, Daniel Duckworth, Alexey Dosovitskiy, Jakob Uszkoreit, Thomas Funkhouser, and Andrea Tagliasacchi. Scene Representation Transformer: Geometry-Free Novel View Synthesis Through Set-Latent Scene Representations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6219–6228, 2022. 1
- [40] Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. Tensor4D: Efficient Neural 4D Decomposition for High-Fidelity Dynamic Reconstruction and Rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16632–16642, 2023. 2
- [41] Liangchen Song, Anpei Chen, Zhong Li, Zhang Chen, Lele Chen, Junsong Yuan, Yi Xu, and Andreas Geiger. NeRF-Player: A Streamable Dynamic Scene Representation with Decomposed Neural Radiance Fields. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2732–2742, 2023. 2
- [42] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Generalizable Patch-Based Neural Rendering. In *European Conference on Computer Vision (ECCV)*, pages 156–174, 2022. 1, 2
- [43] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Light Field Neural Rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8259–8269, 2022.
- [44] Mukund Varma T, Peihao Wang, Xuxi Chen, Tianlong Chen, Subhashini Venugopalan, and Zhangyang Wang. Is Attention All That NeRF Needs? In *International Conference on Learning Representations (ICLR)*, 2023. 1, 2, 3, 6, 12
- [45] Zachary Teed and Jia Deng. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *European Conference on Computer Vision (ECCV)*, pages 402–419, 2020. 2, 5
- [46] Fengrui Tian, Shaoyi Du, and Yueqi Duan. MonoNeRF: Learning a Generalizable Dynamic Radiance Field from Monocular Videos. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17903–17913, 2023. 3
- [47] J. Valmadre and S. Lucey. General trajectory prior for Non-Rigid reconstruction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1394–1401, 2012. 5
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6000–6010, 2017. 2

- [49] Minh Vo, Srinivasa G. Narasimhan, and Yaser Sheikh. Spatiotemporal Bundle Adjustment for Dynamic 3D Reconstruction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1710–1718, 2016. 5
- [50] Feng Wang, Zilong Chen, Guokang Wang, Yafei Song, and Huaping Liu. Masked Space-Time Hash Encoding for Efficient Dynamic Scene Reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, page to appear, 2023. 2
- [51] Feng Wang, Sinan Tan, Xinghang Li, Zeyue Tian, Yafei Song, and Huaping Liu. Mixed Neural Voxels for Fast Multi-view Video Synthesis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19706–19716, 2023. 2
- [52] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yan-shun Zhang, Yingliang Zhang, Minve Wu, Jingyi Yu, and Lan Xu. Fourier PlenOctrees for Dynamic Radiance Field Rendering in Real-time. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13514–13524, 2022. 2, 3
- [53] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snaveley, and Thomas Funkhouser. IBRNet: Learning Multi-View Image-Based Rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4688–4697, 2021. 1, 2, 3
- [54] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4D Gaussian Splatting for Real-Time Dynamic Scene Rendering. *arXiv preprint*, arXiv:2310.08528, 2023. 2
- [55] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time Neural Irradiance Fields for Free-Viewpoint Video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9416–9426, 2021. 2
- [56] 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. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5335–5344, 2020. 2, 3, 6, 7
- [57] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural Radiance Fields from One or Few Images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4576–4585, 2021. 1, 2
- [58] Boyu Zhang, Wenbo Xu, Zheng Zhu, and Guan Huang. Detachable Novel Views Synthesis of Dynamic Scenes Using Distribution-Driven Neural Radiance Fields. *arXiv preprint*, arXiv:2301.00411, 2023. 2
- [59] Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun. NeRF++: Analyzing and Improving Neural Radiance Fields. *arXiv preprint*, arXiv:2010.07492, 2020. 2
- [60] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. 6