

ZeST-NeRF: Using temporal aggregation for Zero-Shot Temporal NeRFs

Violeta Menéndez González^{1,2}
v.menendezgonzalez@surrey.ac.uk

Andrew Gilbert¹
a.gilbert@surrey.ac.uk

Graeme Phillipson²
graeme.phillipson@bbc.co.uk

Stephen Jolly²
stephen.jolly@bbc.co.uk

Simon Hadfield¹
s.hadfield@surrey.ac.uk

¹ Centre for Vision, Speech and Signal
Processing (CVSSP)
University of Surrey
Guildford, UK

² BBC R&D
MediaCityUK
Salford, UK

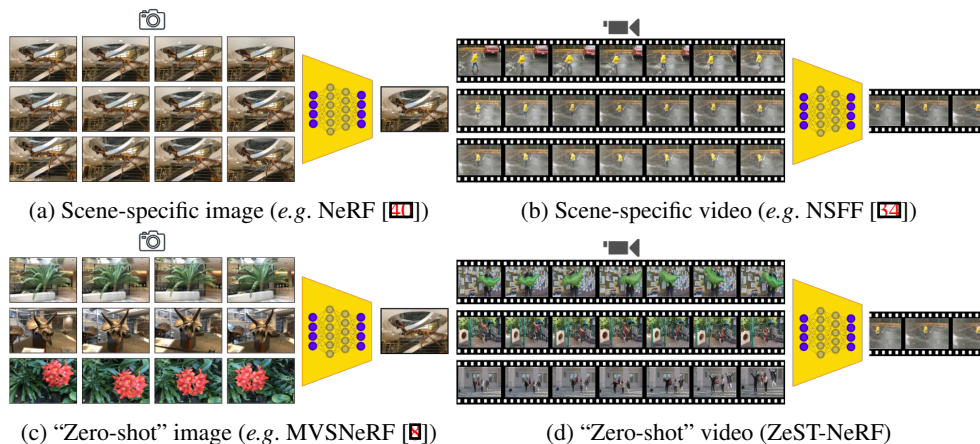


Figure 1: “Zero-shot” temporal NeRF contrasted with existing techniques

Abstract

In the field of media production, video editing techniques play a pivotal role. Recent approaches have had great success at performing novel view image synthesis of static scenes. But adding temporal information adds an extra layer of complexity. Previous models have focused on implicitly representing static and dynamic scenes using NeRF. These models achieve impressive results but are costly at training and inference time. They overfit an MLP to describe the scene implicitly as a function of position. This paper proposes ZeST-NeRF, a new approach that can produce temporal NeRFs for new scenes without retraining. We can accurately reconstruct novel views using multi-view synthesis techniques and scene flow-field estimation, trained only with unrelated scenes. We demonstrate how existing state-of-the-art approaches from a range of fields cannot adequately solve this new task and demonstrate the efficacy of our solution. The resulting network improves quantitatively by 15% and produces significantly better visual results.

1 Introduction

Producing photo-realistic renderings of natural scenes under new viewpoints is important for media content production, virtual/augmented reality, and image/video editing. The challenge increases significantly when the scene is time-dependent, i.e. when the camera or subjects undergo movement.

In the area of novel view synthesis, many new approaches [8, 12] rely on the popular implicit representation given by Neural Radiance Fields (NeRF) [40] (Fig. 1a). This has been extended to dynamic scene methods that aim to model the temporal dimension alongside the spatial ones [53, 54] (Fig. 1b). These approaches achieve impressive photo-realistic results but have similar limitations to NeRF, being very expensive, requiring a lot of input views, and having a very long per-scene optimisation process. Some approaches [8, 19, 71] worked towards generalising NeRFs to unseen scenes and reducing the required number of input views (Fig. 1c). This is already a highly ill-defined problem, which suffers from uncertainty and ambiguity in the reconstructed data. Naively adding a layer of temporal complexity can significantly exacerbate this. In this work, the temporal data becomes an advantage, as information is aggregated across frames to resolve ambiguities in unknown scenes.

In particular, we propose ZeST-NeRF (Fig. 1d), which uses temporal information to reconstruct a scene geometry and motion estimate. This can then be used to generalise to unseen scenes without requiring expensive retraining. Given a set of keyframes, we use a static geometry encoding volume to inform both a static NeRF (which reconstructs the background) and a dynamic encoding volume (which reconstructs the motion between frames). We then combine these to generate a new video from a new point of view. Because our approach learns to estimate structure and motion in a scene-agnostic way, it can be applied “zero-shot”¹ to new scenes without laborious training. It uses temporal and spatial information to recover missing information in unseen areas. It can be further fine-tuned for short amounts of time to increase the quality of the reconstructions. We will make our code publicly available². In summary, the contributions of this paper are:

- First radiance field approach capable of “zero-shot” novel viewpoint rendering in completely unseen videos of complex scenes.
- An efficient multi-encoding-volume approach to scene-agnostic video representation.
- A new evaluation protocol and newly developed baselines for the problem of “zero-shot” novel-viewpoint-rendering in the video.

2 Background

Novel view synthesis Classic Image-based rendering (IBR) techniques [0, 12, 59] usually attempt to model an intermediate scene geometry. These geometrical representations are based on restrictive structures like voxel grids [62, 63], point clouds [62, 67], or multi-layer arrangements [16, 17, 58, 76]. In recent years, implicit representations have been prevalent when approaching novel view synthesis techniques. NeRF [40] proposed an entirely neural scene representation. A Multi-Layer Perceptron (MLP) is used to parameterise a function rendering density and colour by querying 3D location and viewing direction. This approach revolutionised the field; however, in its original form, NeRF is very costly to run. In addition, this model has a long per-scene optimisation process, which prevents it from being useful in many important applications. Following approaches worked towards improving performance

¹Note that in the context of NeRFs we use ‘zero-shot’ to refer to the ability to perform inference on scenes (and scene configurations) that differ from those seen during training. This is in slight contrast to the usage of the phrase in supervised learning which refers specifically to ‘categories’ of data which were unseen during training

²<https://github.com/violetamenendez/zest-nerf>

and increasing flexibility [10, 11, 88, 57]. Regardless, these models require dense input views, are costly, and don't generalise to new scenes.

Some recent approaches have relaxed the requirement for a high number of input views by applying data augmentation [9] or by introducing regularisations [13, 29, 44, 50]. Furthermore, some approaches attempt to generalise their models to apply to new scenes that the models haven't been trained on [27, 70]. This is achieved by introducing neural geometry priors based on IBR approaches, like 2D feature aggregation [63], stereo-matching techniques [10], or recurrent aggregation [73]. MVNeRF [8] constructs a 3D feature encoding volume based on Plane Sweep Volumes [16]. The model only needs three input images at training time and can still generalise to new scenes. Nevertheless, these models suffer from significant artefacts and incorrect outputs in occluded areas that are not visible from the reference view. SVS [49] tries to solve this by applying adversarial training to hallucinate content in the regions occluded in all inputs. While this improves results in many cases, the adversarial regime harms training stability. Our paper extends this 'encoding volume' approach to "zero-shot" transfer, developing both a static and dynamic encoding volume for application to temporal sequences.

Dynamic scene representation Within dynamic scene reconstruction, many approaches require ground truth RGBD or costly hand-labelled data. Some methods concentrate on recovering either a single non-rigid 3D object [8, 12, 22, 24, 32, 43, 48, 74, 78], or sparse geometry and motion [45, 55, 62, 75]. Yet others leverage monocular depth estimation models and image segmentation approaches to decompose scenes into rigid and non-rigid areas [61, 46, 49, 64, 69]. Some transform a single canonical radiance field with rigid transformations or modelled dynamics [15, 46, 47, 48, 60, 72]. And others attempt to use other pre-computed scene representations to improve rendering efficiency [85, 66]. Recently some techniques have explored estimating scene flow [20] for this [6, 23, 26, 47, 41, 68].

When reconstructing scene geometry for novel view synthesis, many approaches require multiple input views from synchronised videos [11, 9, 77]. Yoon *et al.* [70] suggest a self-supervised depth fusion network DFNet. This model combines dense view-dependent monocular depth and sparse view-independent multi-view stereo to explain dynamic scene geometry. However, this approach required synchronised annotated data and produced significant artefacts in disocclusions. Instead, other approaches [18, 54] propose a space-time synthesis model that extends NeRF [40] to estimate scene flow fields. They deal with occlusions by predicting occlusion weights or having 3D flow supervision. Still, they suffer from similar computational limitations as other NeRF models. Li *et al.* [63] use importance sampling and hierarchical training to boost training speed. Nevertheless, all these methods require per-scene optimisation processes. In contrast, we aim to render new scenes the model hasn't trained on in a "zero-shot" manner.

3 Approach

3.1 Geometry and motion encoding volumes

Our architecture (See Fig. 2) exploits the power of multi-view 3D CNN encoding volumes [8, 19] to generalise to new scenes. This 3D volume integrates 2D CNN features of the input images by warping multiple sweeping planes of source view features. This differs from techniques like Deep Stereo [16], which perform plane sweeps using the raw colour pixels to produce their correlation volume. This allows us to generalise correlations between images, which can then be used to reason about geometry and motion, helping the network generalise to previously unseen scenes.

We use two different volumes which learn different correlations between frames by construction. The *geometry volume* and the *motion volume* described below are constructed with the same network architecture but have different inputs and do not share weights.

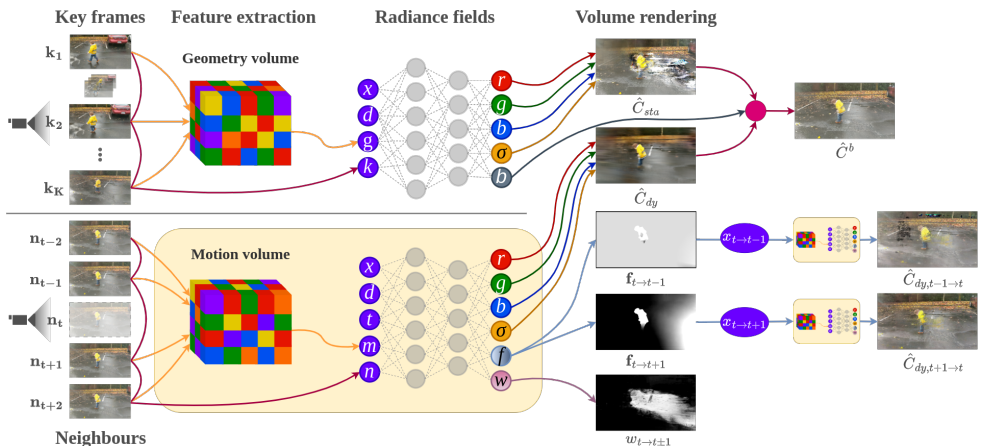


Figure 2: **Model overview:** ZeST-NeRF uses a geometry and a motion encoding volume to inform the static and the dynamic NeRFs, respectively.

Geometry volume Given the image sequence $\mathcal{V} = \{\mathbf{I}_i \mid \mathbf{I}_i \in \mathbb{R}^{H \times W \times 3}\}_{i=1}^N$, we use K key-frames that are representative of the whole video $\mathcal{K} = \{\mathbf{I}_i \mid \mathbf{I}_i \in \mathcal{V}, i = \lambda \cdot j, \lambda \in \mathbb{N}\}_{j=1}^K$. In our experiments, we use key-frames equally spread across the sequence. However, our approach could easily make use of intelligent key-frame selection techniques. These frames will inform the volume about the general geometry of the scene, particularly including static background elements. We then extract the deep features of each key frame using a deep 2D convolutional network $E(\mathbf{I}_i | \mathbf{w}_\Psi)$ with weights \mathbf{w}_Ψ . This network consists of downsampling convolutional layers, batch-normalisation and ReLU activation layers.

We build the plane sweep volume by aligning each feature map to the reference view at multiple depths. To achieve this, a homography $\mathcal{H}_i(d)$ is computed for each view at each depth. Given the camera parameters $\{\mathbf{K}_i, \mathbf{R}_i, \mathbf{t}_i\}$ (intrinsic, rotation and translation) for camera i the homography is defined as

$$\mathcal{H}_i(d) = \mathbf{K}_i \cdot \mathbf{R}_i \cdot \left(\mathbb{I} + \frac{(\mathbf{t}_{ref} - \mathbf{t}_i) \cdot \mathbf{n}_{ref}^T}{d} \right) \cdot \mathbf{R}_{ref}^T \cdot \mathbf{K}_{ref}^T \quad (1)$$

where \mathbb{I} is the 3×3 identity matrix, \mathbf{n}_{ref} the principle axis of the reference camera, and d is the depth to which the images are being warped. This operation is differentiable, which allows for end-to-end training of the feature encoding network weights \mathbf{w}_Ψ based on the downstream reconstruction losses.

The feature sweep volumes $\mathcal{S}(\mathbf{I}_i) = \{W(\mathcal{H}_i(d), E(\mathbf{I}_i | \mathbf{w}_\Psi)) \mid \forall d \in \{1, \dots, D\}\}$ for each keyframe $\mathbf{I}_i \in \mathcal{K}$ are created by applying the warping function W determined by the homography \mathcal{H} to the keyframe’s feature maps at every depth. Then, a variance-based cost volume $\mathbf{C}(\mathcal{K}) = \text{Var}(\{\mathcal{S}(\mathbf{I}_i) \mid \forall \mathbf{I}_i \in \mathcal{K}\})$ [10, 59] is generated by aggregating all the warped feature sweep volumes. This cost volume encodes appearance variations across views. A variance based metric makes it possible to compute this using an arbitrary number of input views.

Finally, the cost volume is processed using a 3D CNN UNet-like network [52]. The output of this network is the neural geometry volume $\mathbf{G} = V(\mathbf{C}(\mathcal{K}) | \mathbf{w}_\Omega)$ (Fig. 3).

$$\mathbf{G} = V(\mathbf{C}(\mathcal{K}) | \mathbf{w}_\Omega) \quad (2)$$

This embedding volume encodes the feature correlations across key views in the video, which have been propagated using downsampling and upsampling layers with skip connections. This

structure allows the volume to represent the static appearance elements from the video. In this way, the system can generalise to new scene arrangements and video lengths at inference time.

Motion volume This volume is similar in concept to the *geometry volume*, but it extracts dynamic correlations across short-term neighbouring frames instead of modelling static structure. By choosing only frames near our target time, the network is informed of the correlations caused by moving objects (Fig. 3).

We choose M neighbours around our target frame at time t to create this motion encoding volume. In practice, we use $M = 4$ for neighbours $\mathcal{N}_t = \{\mathbf{I}_i \mid \mathbf{I}_i \in \mathcal{V}, i \in \{t \pm 1, t \pm 2\}\}$. We build the motion volume following the same approach as for the static geometry encoding volume. Using the neighbouring frames $\mathbf{I}_i \in \mathcal{N}_t$ and the homography in Equation 1 to build the feature sweep volumes $S(\mathbf{I}_i)$. Then aggregating the feature volumes in the cost volume $\mathbf{C}(\mathcal{N}_t)$. And subsequently, passing the cost volume through a 3D CNN network with the same architecture as for the *geometry volume*, but with different training weights. We then obtain the motion volume,

$$\mathbf{M} = V(\mathbf{C}(\mathcal{N}_t) \mid \mathbf{w}_{\Xi}) \quad (3)$$

This volume thus represents the short-term behaviours of dynamic scene elements.

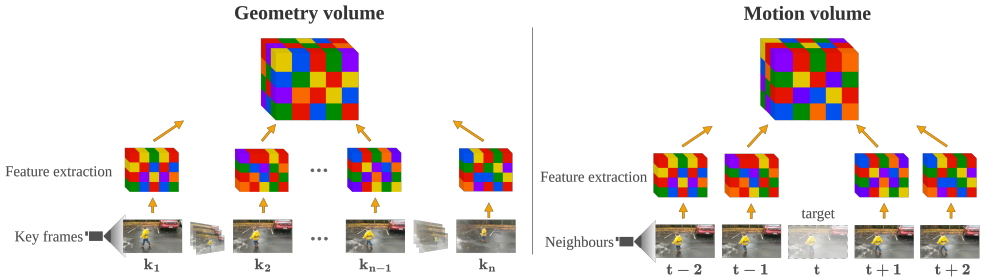


Figure 3: **Volumes:** Static (left) and Dynamic (right) encoding volumes

3.2 Static neural radiance fields

We optimise an MLP [40] F_{Θ} with parameters \mathbf{w}_{Θ} to decode the *geometry volume* embedding into a density and view-dependent radiance (colour). Given a 3D point x and a viewing direction d , the network F_{Θ} regresses the density σ and colour c at that point, conditioned on the *geometry volume* \mathbf{G} from Equation 2. To allow the correlations and structures in \mathbf{G} to be mapped back to the original scene albedo, we use the pixel colour of the original key frame inputs \mathcal{K} as additional conditioning information [8]. We also predict blending weights b [52], which assign the linear weights to blend the colour and density estimated by the static and dynamic radiance fields (see Sec. 3.4). These weights are unsupervised and give higher importance to the static regions of the scene that this static representation can best model.

$$F_{\Theta}: (x, d, \mathbf{G}, \mathcal{K} \mid \mathbf{w}_{\Theta}) \mapsto (\sigma_x, c_{x,d}, b_{x,d}) \quad (4)$$

3.3 Dynamic neural radiance fields

To model the dynamics of our scene, we cannot just naively add a time dimension to our radiance field input. Doing so results in very noisy and inconsistent results due to the problem’s dimensionality [63, 64]. Instead, we estimate scene flow fields [64] to aggregate information between frames. Our dynamic representation predicts the forward and backward 3D scene flow at a given location $\mathcal{F}_t = (\mathbf{f}_{t \rightarrow t+1}, \mathbf{f}_{t \rightarrow t-1})$. This denotes 3D offset vectors that point to the position of that point at times $t + 1$ and $t - 1$.

Unfortunately, density and scene flow are ambiguous in disocclusion regions caused by 3D motion. There is no simple mechanism to distinguish a region with a low correlation due to being empty and one with a low correlation due to being occluded in all frames. The prevalence of these fully occluded regions grows drastically as the number of input frames is reduced, leading traditional radiance field models to produce reconstructions full of unrealistic holes. To avoid this, we also predict disocclusion weights $\mathcal{W}_t = (w_{t \rightarrow t+1}, w_{t \rightarrow t-1})$. These unsupervised weights can be seen as the confidence of the estimated results and can guide the application strength of the reconstruction losses to the areas we are certain are observable across the video. The dynamic representation is then given by the following function F_Φ where x is the query point in 3D space, with viewing direction d at time $t \in \{1, \dots, N\}$. \mathbf{M} is our motion encoding volume from Equation 3, to which we concatenate the original colour values from neighbouring frames \mathcal{N}_t .

$$F_\Phi: (x, d, t, \mathbf{M}, \mathcal{N}_t | \mathbf{w}_\Phi) \mapsto (\sigma_{x,t}, c_{x,d,t}, \mathcal{F}_t, \mathcal{W}_t). \quad (5)$$

3.4 Volume rendering

We use differentiable ray marching to render the colour of image pixels. This is done by projecting (“marching”) a ray \mathbf{r}_t through a pixel in the target image \mathbf{I}_t at time t . We query the respective neural radiance networks at regular intervals $\gamma \in [1.. \text{inf}]$ along this ray to obtain the radiance (colour) $c_\gamma = c(\gamma)$ and density $\sigma_\gamma = \sigma(\gamma)$ at each ray sample.

$$(\sigma_\gamma, c_\gamma, b_\gamma) = F_\Theta(\mathbf{r}_{t,\gamma}, \mathbf{d}, \mathbf{G}, \mathcal{K} | \mathbf{w}_\Theta), \quad (\sigma_{t,\gamma}, c_{t,\gamma}, \mathcal{F}_{t,\gamma}, \mathcal{W}_{t,\gamma}) = F_\Phi(\mathbf{r}_{t,\gamma}, \mathbf{d}, t, \mathbf{M}, \mathcal{N}_t | \mathbf{w}_\Phi) \quad (6)$$

where $\mathbf{r}_{t,\gamma} = \mathbf{r}_t(\gamma) = \mathbf{o}_t + \gamma \mathbf{d}$ with \mathbf{o}_t being the ray origin and \mathbf{d} the direction vector. Then, we use the predicted blending weights b_γ to “blend” the colour and density samples from F_Θ and F_Φ . Obtaining the estimated colour $\hat{C}^b(\mathbf{r}_t)$ of the pixel via the volume rendering Equation [23]:

$$\hat{C}_{sta}^b(\mathbf{r}_t) = \sum_\gamma \tau_\gamma^b (1 - \exp(-\sigma_\gamma)) c_\gamma, \quad \hat{C}_{dy}^b(\mathbf{r}_t) = \sum_\gamma \tau_\gamma^b (1 - \exp(-\sigma_{t,\gamma})) c_{t,\gamma} \quad (7)$$

$$\hat{C}^b(\mathbf{r}_t) = (1 - b_\gamma) \hat{C}_{sta}^b(\mathbf{r}_t) + b_\gamma \hat{C}_{dy}^b(\mathbf{r}_t) \quad (8)$$

where τ_γ^b is the blended transmittance at sample γ , which represents the probability that the ray travels up to γ without hitting another particle.

$$\tau_\gamma^b = \exp\left(-\sum_{j=1}^{\gamma-1} (\sigma_j (1 - b_j) + \sigma_{t,j} b_j)\right) \quad (9)$$

3.5 Losses and regularisations

Reconstruction loss To enforce low-frequency correctness in the output, we compute the L_2 loss between the blended colour estimate and the true colour,

$$\mathcal{L}_{rec} = \|\hat{C}^b(\mathbf{r}_t) - C(\mathbf{r}_t)\|_2^2 \quad (10)$$

Temporal photometric consistency To encourage the reconstruction of the scene at time t to be consistent with the scene at neighbouring times $k \in \mathcal{N}(t)$. We apply a photometric loss [54] to our dynamic network, which considers the motion due to 3D scene flow. This is done by warping the scene from neighbouring frame k to time t and comparing it to the ground truth scene at time t . Because this loss suffers from ambiguity in disocclusion areas, it is scaled by the estimated confidence weights mentioned in Sec. 3.3.

$$\mathcal{L}_{pho} = \sum_{k \in \mathcal{N}(t)} \hat{W}_{k \rightarrow t}(\mathbf{r}_t) \|\hat{C}_{dy}(\mathbf{r}_t + \mathbf{f}_{t \rightarrow k}) - C(\mathbf{r}_t)\|_2^2 \quad (11)$$

where $\mathbf{f}_{t \rightarrow k} \in \mathcal{F}_t$ is the scene flow field from time t to k . We apply volume rendering to get the colour $\hat{C}_{dy}(\mathbf{r}_t + \mathbf{f}_{t \rightarrow k})$ from the dynamic network F_Φ at time k . We also apply volume rendering to get the disocclusion weights $\hat{W}_{k \rightarrow t}(\mathbf{r}_t)$, which are the accumulated confidence estimates $w_{t \rightarrow k, \gamma}$ along each ray.

$$\hat{C}_{dy}(\mathbf{r}_t + \mathbf{f}_{t \rightarrow k}) = \sum_{\gamma} \tau_{k, \gamma} (1 - \exp(-\sigma_{k, \gamma})) c_{k, \gamma}, \quad \tau_{k, \gamma} = \exp\left(-\sum_{j=1}^{\gamma-1} (\sigma_{k, j})\right) \quad (12)$$

$$\hat{W}_{k \rightarrow t}(\mathbf{r}_t) = \sum_{\gamma} \tau_{k, \gamma} (1 - \exp(-\sigma_{k, \gamma})) w_{t \rightarrow k, \gamma} \quad (13)$$

Disocclusion weight regularisation The system can fall into a degenerate local minimum, where the disocclusion weights $w_{t \rightarrow k}$ are zero at every 3D point x_t , and only the static NeRF is active. We avoid this with an L_1 regularisation that pushes them closer to one.

$$\mathcal{L}_w = \sum_{x_t} \sum_{k \in \mathcal{N}(t)} \|w_{t \rightarrow k}(x_t) - 1\| \quad (14)$$

Blending weights entropy loss This loss encourages blending weight to be either 0 or 1, which can help to reduce the ghosting caused by learned semi-transparent blending weights.

$$\mathcal{L}_b = \|-b \cdot \log(b)\| \quad (15)$$

Cycle loss We enforce a cyclic regularisation to ensure that the predicted forward flow field at time t ($\mathbf{f}_{t \rightarrow k}$) is consistent with the backward flow field at time k ($\mathbf{f}_{k \rightarrow t}$), where k is a neighbouring time $k \in \{t \pm 1\}$. Fundamentally, if $x_{t \rightarrow k} = x_t + \mathbf{f}_{t \rightarrow k}$ then we should see that $x_{t \rightarrow k} + \mathbf{f}_{k \rightarrow t} = x_t$. That is, $\mathbf{f}_{t \rightarrow k} = -\mathbf{f}_{k \rightarrow t}$. Flow fields are ambiguous at disocclusion areas, so we regulate this loss with the predicted disocclusion weights [54].

$$\mathcal{L}_{cyc} = \sum_{x_t} \sum_{k \in \{t \pm 1\}} w_{t \rightarrow k} \|\mathbf{f}_{t \rightarrow k}(x_t) + \mathbf{f}_{k \rightarrow t}(x_{t \rightarrow k})\| \quad (16)$$

Scene flow regularisation We assume that motion is small in most 3D space [61], so we minimise the absolute value of the flow fields.

$$\mathcal{L}_{min} = \sum_{x_t} \sum_{k \in \{t \pm 1\}} \|\mathbf{f}_{t \rightarrow k}(x_t)\| \quad (17)$$

Scene flow spatial smoothness We assume that the scene deforms in a piece-wise smooth way [43], meaning the scene flow field is spatially smooth.

$$\mathcal{L}_{sp} = \sum_{x_t} \sum_{y_t \in \mathcal{N}(x_t)} \sum_{k \in \{t \pm 1\}} w^{dist}(x_t, y_t) \|\mathbf{f}_{t \rightarrow k}(x_t) - \mathbf{f}_{t \rightarrow k}(y_t)\| \quad (18)$$

where $\mathcal{N}(x_t)$ are the neighbouring points of x_t , and $w^{dist}(x_t, y_t) = \exp(-2\|x_t - y_t\|)$ are weightings based on Euclidean distance.

Scene flow temporal smoothness Finally, we add a scene flow temporal smoothness loss. This loss encourages 3D point trajectories to have minimal kinetic energy [62] i.e. constant velocity and piece-wise linear motion.

$$\mathcal{L}_{temp} = \sum_{x_t} \|\mathbf{f}_{t \rightarrow t+1}(x_t) + \mathbf{f}_{t \rightarrow t-1}(x_t)\|_2^2 \quad (19)$$

3.6 Weakly-supervised pre-training

Since the problem of reconstructing complex dynamic scenes is extremely ill-posed, the losses can converge to local minima when randomly initialised [54]. We first complete a data-mining stage to initialise the problem optimally. We use monocular optical flow and depth estimation networks to generate pseudo ground truth data to guide our network. However, as both models are not completely accurate, we use them for initialisation only and decay their contribution to zero during training.

Geometric consistency First, we compute a reprojection error between the scene flow field estimated by ZeST-NeRF and that derived from pre-trained optical flow models [59]. Given a ray \mathbf{r}_t at time t and estimated 3D scene-flow $\mathbf{f}_{t \rightarrow k}$, we can calculate the expected location of the 3D point $\hat{X}_t(\mathbf{r}_t)$ and the expected 2D optical flow $\hat{F}_{t \rightarrow k}(\mathbf{r}_t)$ of that point, by using volume rendering,

$$\hat{X}_t(\mathbf{r}_t) = \sum_{\gamma} \tau_{t,\gamma} (1 - \exp(-\sigma_{t,\gamma})) x_{t,\gamma}, \quad \hat{F}_{t \rightarrow k}(\mathbf{r}_t) = \sum_{\gamma} \tau_{t,\gamma} (1 - \exp(-\sigma_{t,\gamma})) \mathbf{f}_{t \rightarrow k,\gamma} \quad (20)$$

Then we project that displaced 3D point $\hat{X}_t(\mathbf{r}_t) + \hat{F}_{t \rightarrow k}(\mathbf{r}_t)$ into the camera view of the neighbouring frame (k) as $\hat{p}_{t \rightarrow k} = \pi(\mathbf{K}(\mathbf{R}_k(\hat{X}_t(\mathbf{r}_t) + \hat{F}_{t \rightarrow k}(\mathbf{r}_t)) + \mathbf{t}_k))$. We compare our estimation to the displaced pixel $p_{t \rightarrow k} = p_t + \mathbf{u}_{t \rightarrow k}$, where $\mathbf{u}_{t \rightarrow k}$ is the 2D optical flow derived using a pre-trained optical flow model [59]. We can then minimise the reprojection error with the following loss [45],

$$\mathcal{L}_{geo} = \sum_{k \in \{t \pm 1\}} \|\hat{p}_{t \rightarrow k} - p_{t \rightarrow k}\| \quad (21)$$

Single-view depth prior In traditional *neural radiance field* approaches, depth maps are extremely noisy, if not nonsensical. We use a pre-trained monocular depth estimation network [60] to encourage the expected termination depth along each ray $\hat{D}(\mathbf{r}_t)$ to be close to the derived depth $D(\mathbf{r}_t)$. We apply an L_1 loss using a robust scale-shift invariant metric [32],

$$\mathcal{L}_{depth} = \|\hat{D}(\mathbf{r}_t) - D(\mathbf{r}_t)\|. \quad (22)$$

4 Evaluation

Implementation details We use the same architecture to build both the *geometry* and *motion* volumes, but they are optimised separately and thus do not share weights. We extract 32 feature channels from the input images using a 2D CNN consisting of downsampling convolutional layers with Batch-Normalisation (BN) and ReLU activation function. Warping these features we create a Plane Sweep Volume per input image, selecting $D = 128$ depth planes. These are aggregated through a variance-based cost volume and processed using a 3D UNet-like network, consisting of downsampling and upsampling Convolutional layers with BN and ReLU, and skip connections.

For the NeRF MLPs, we follow a similar setup to the original case [40]. We sample 128 points along each ray, with a ray batch of 1024. We also have two separate networks for the static and dynamic parts, which do not share weights. We append the normalised time indices in NSFF [54] to our dynamic network inputs. The MLP networks return the estimated colour c and density σ , as well as blending weights b in the case of the Static MLP, and 3D scene flow \mathbf{f} and occlusion weights w in the case of the Dynamic MLP. We use an Adam optimiser [60] with a learning rate of $5e - 4$. We use positional encoding (PE) [40] for the 3D location and viewing direction before feeding them into the networks. For more detailed information about the architecture, refer to the supplementary material.

Table 1: **Quantitative evaluation.** Results on cross-validation training and further fine-tuning. **Bold** is best result, *italic* is second best.

| Model | Scene-agnostic | | | Scene-specific | | |
|------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
| | PSNR \uparrow | SSIM \uparrow | LPIPS \downarrow | PSNR \uparrow | SSIM \uparrow | LPIPS \downarrow |
| NSFF [64] | 13.15 | 0.3178 | 0.7334 | 28.19 | 0.9280 | 0.0450 |
| SVS [19] | 17.91 | 0.4958 | 0.3698 | 21.28 | 0.7103 | 0.2001 |
| MVSNeRF [8] | <i>19.37</i> | <i>0.6198</i> | <i>0.2885</i> | <i>26.94</i> | <i>0.8678</i> | 0.1208 |
| ZeST-NeRF (Ours) | 21.59 | 0.6239 | 0.2048 | 26.94 | 0.8575 | <i>0.0995</i> |

Baseline models and datasets We compare our results to a space-time view synthesis approach NSFF [64] and two scene-agnostic static view synthesis approaches SVS [19] and MVSNeRF [8] applied naively across frames. We attempted to reproduce the results in Li *et al.* [63] but found it impossible with the information given in the paper due to the lack of published code. We train on the Dynamic Scenes dataset [70]. This dataset consists of 8 short-time video clips recorded with 12 synchronised cameras, and sampled from each camera at different times to simulate a moving monocular camera. The scenes are collected in the wild and feature complex motions such as jumping, running, or dancing. As this dataset is quite small, we perform a leave-one-out cross-validation study to prove our performance. We train one full model per subset of scenes, created by holding out one single scene for validation. This means that we can test on every scene of the Dynamic Scenes dataset with a model that has been trained on completely unrelated scenes. Please refer to the supplementary material for additional details regarding the experimental setup.



Figure 4: **Qualitative results:** Example results (see supplementary material for more).

Results When comparing ZeST-NeRF to other state-of-the-art models in Table 1, we can see that our approach outperforms all models when applied in a scene-agnostic manner. ZeST-NeRF improves the results of other scene-agnostic methods by at least 15% across all metrics. In addition, a ZeST-NeRF model pre-trained on different sequences can be rapidly fine-tuned on a specific scene to produce results competitive with exhaustively optimised scene-specific approaches quickly. Qualitative results are shown in Figure 4 (see supplementary material for additional qualitative examples). We can observe that our model produces significantly more accurate and pleasant results than the other methods. Approaches like MVSNerF [8] or SVS [19] are able to reconstruct the correct background, but struggle to recover the dynamic objects in the scene. As the numbers from the ablation study suggest (Table 2), this is probably due to the necessity of having both a Geometry and a Dynamic network to recover static and dynamic features. On the other hand, the NSFF [32] model is not really equipped to solve this scene-agnostic problem.

Despite the inherent complexity and ill-posed nature of the scene-agnostic “zero-shot” novel view synthesis problem, it is important to acknowledge that our approach is not without its limitations. As the network concentrates on recovering dynamic areas, the background reconstruction suffers from a slight quality reduction and some blurriness may persist in certain scenes. Furthermore, our dynamic reconstruction occasionally suffers from duplication artefacts. Despite this, our approach is a significant step-change in performance compared to previous research.

Table 2: **Ablation study.** Effect of each addition to the model. **Bold** is best result.

| Geometry | Dynamic | PSNR \uparrow | SSIM \uparrow | LPIPS \downarrow |
|----------|---------|-----------------|-----------------|--------------------|
| × | × | 12.67 | 0.4142 | 0.8277 |
| ✓ | × | 17.87 | 0.4533 | 0.4781 |
| × | ✓ | 20.19 | 0.5814 | 0.2447 |
| ✓ | ✓ | 21.59 | 0.6239 | 0.2048 |

5 Conclusions

In this paper we have proposed ZeST-NeRF. This dynamic-scene representation technique combines multi-view synthesis and scene flow field estimation approaches to rendering novel views in previously unseen scenes. This is useful for applications in media production, where we have many videos and want to generate new camera views without expensive scene-specific training.

Attempting to apply previous state-of-the-art techniques from related fields naively, leads to blurry results with overwhelming numbers of artefacts or, in some cases, a complete inability to train. The problem of scene-agnostic “zero-shot” novel view synthesis is highly ill-posed, and our model exhibits certain limitations. Artefacts like duplications persist in certain scenes, particularly in regions characterised by temporal motion. In addition, some blurriness may occur in background areas. Furthermore, given the limited training dataset, our generalisation power may be constrained by the type of scene dynamics and length of motion. It would be interesting to study the behaviour of the model in larger datasets with a more varied set of motions and scene complexities. Nonetheless, our approach constitutes a significant advancement in performance compared to prior research efforts.

In the future, it would be interesting to explore techniques that can better reconstruct dynamic objects, potentially aided by stereo views from multiple cameras. Other representation techniques, like instant Neural Graphic Primitives [42], may help improve efficiency. In addition, generative models like diffusion models could be explored to tackle uncertainty in highly occluded areas or areas with more significant motion.

References

- [1] 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 (CVPR)*, 2020.
- [2] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [4] Mojtaba Bermana, K. Myszkowski, H. Seidel, and Tobias Ritschel. X-Fields: Implicit Neural View-, Light- and Time-Image Interpolation. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2020)*, October 2020.
- [5] Aljaž Božič, Michael Zollhöfer, Christian Theobalt, and Matthias Nießner. DeepDeform: Learning Non-Rigid RGB-D Reconstruction With Semi-Supervised Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [6] Fabian Brickwedde, Steffen Abraham, and Rudolf Mester. Mono-SF: Multi-View Geometry Meets Single-View Depth for Monocular Scene Flow Estimation of Dynamic Traffic Scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [7] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM transactions on graphics (TOG)*, June 2013.
- [8] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [9] Tianlong Chen, Peihao Wang, Zhiwen Fan, and Zhangyang Wang. Aug-NeRF: Training Stronger Neural Radiance Fields With Triple-Level Physically-Grounded Augmentations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [10] Shuo Cheng, Zexiang Xu, Shilin Zhu, Zhuwen Li, Li Erran Li, Ravi Ramamoorthi, and Hao Su. Deep Stereo Using Adaptive Thin Volume Representation With Uncertainty Awareness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [11] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo Radiance Fields (SRF): Learning View Synthesis for Sparse Views of Novel Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.

- [12] Paul Debevec, Yizhou Yu, and George Borshukov. Efficient View-Dependent Image-Based Rendering with Projective Texture-Mapping. In *Rendering Techniques '98. EGSR 1998. Eurographics*. 1998.
- [13] Yu Deng, Jialong Yang, Jianfeng Xiang, and Xin Tong. GRAM: Generative Radiance Manifolds for 3D-Aware Image Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [14] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, Pushmeet Kohli, Vladimir Tankovich, and Shahram Izadi. Fusion4D: Real-time performance capture of challenging scenes. *ACM transactions on graphics (TOG)*, July 2016.
- [15] Yilun Du, Yanan Zhang, Hong-Xing Yu, Joshua B. Tenenbaum, and Jiajun Wu. Neural Radiance Flow for 4D View Synthesis and Video Processing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, September 2021.
- [16] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. DeepStereo: Learning to Predict New Views from the World’s Imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [17] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [18] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic View Synthesis from Dynamic Monocular Video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, May 2021.
- [19] Violeta Menéndez González, Andrew Gilbert, Graeme Phillipson, Stephen Jolly, and Simon Hadfield. SVS: Adversarial refinement for sparse novel view synthesis. In *Proceedings of the 33rd British Machine Vision Conference (BMVC)*, 2022.
- [20] Simon Hadfield and Richard Bowden. Scene particles: Unregularized particle-based scene flow estimation. *IEEE transactions on pattern analysis and machine intelligence*, 2013.
- [21] Tao Hu, Shu Liu, Yilun Chen, Tiancheng Shen, and Jiaya Jia. EfficientNeRF Efficient Neural Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [22] Zeng Huang, Tianye Li, Weikai Chen, Yajie Zhao, Jun Xing, Chloe LeGendre, Linjie Luo, Chongyang Ma, and Hao Li. Deep Volumetric Video From Very Sparse Multi-view Performance Capture. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [23] Junhwa Hur and Stefan Roth. Self-Supervised Monocular Scene Flow Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [24] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. VolumeDeform: Real-Time Volumetric Non-rigid Reconstruction. *12th European Conference on Computer Vision (ECCV)*, 2016.

- [25] Matthias Innmann, Kihwan Kim, Jinwei Gu, Matthias Nießner, Charles Loop, Marc Stamminger, and Jan Kautz. NRMVS: Non-Rigid Multi-View Stereo. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, March 2020.
- [26] Huaizu Jiang, Deqing Sun, Varun Jampani, Zhaoyang Lv, Erik Learned-Miller, and Jan Kautz. SENSE: A Shared Encoder Network for Scene-Flow Estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [27] Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. GeoNeRF: Generalizing NeRF With Geometry Priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [28] James T. Kajiya and Brian P. Von Herzen. Ray Tracing Volume Densities. *ACM SIGGRAPH*, 1984.
- [29] Mijeong Kim, Seonguk Seo, and Bohyung Han. InfoNeRF: Ray Entropy Minimization for Few-Shot Neural Volume Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.
- [30] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations (ICLR)*, 2014.
- [31] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Monocular Dense 3D Reconstruction of a Complex Dynamic Scene from Two Perspective Frames. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2017.
- [32] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural Human Performer: Learning Generalizable Radiance Fields for Human Performance Rendering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [33] Tianye Li, Mira Slavcheva, Michael Zollhöfer, 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [34] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural Scene Flow Fields for Space-Time View Synthesis of Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [35] Kai-En Lin, Lei Xiao, Feng Liu, Guwei Yang, and Ravi Ramamoorthi. Deep 3D Mask Volume for View Synthesis of Dynamic Scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021.
- [36] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, August 2020.
- [37] Zhaoyang Lv, Kihwan Kim, Alejandro Troccoli, Deqing Sun, James M. Rehg, and Jan Kautz. Learning rigidity in dynamic scenes with a moving camera for 3d motion field estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [38] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

- [39] Leonard McMillan and Gary Bishop. Plenoptic modeling: An image-based rendering system. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1995.
- [40] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [41] Himangi Mittal, Brian Okorn, and David Held. Just Go With the Flow: Self-Supervised Scene Flow Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [42] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM transactions on graphics (TOG)*, July 2022.
- [43] Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [44] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. RegNeRF: Regularizing Neural Radiance Fields for View Synthesis From Sparse Inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [45] Hyun Soo Park, Takaaki Shiratori, Iain Matthews, and Yaser Sheikh. 3D reconstruction of a moving point from a series of 2D projections. In *Proceedings of the 11th European Conference on Computer Vision (ECCV)*, September 2010.
- [46] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable Neural Radiance Fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021.
- [47] 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 (TOG)*, December 2021.
- [48] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [49] René Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun. Dense Monocular Depth Estimation in Complex Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [50] 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*, March 2022.
- [51] Daniel Rebain, Mark Matthews, Kwang Moo Yi, Dmitry Lagun, and Andrea Tagliasacchi. LOLNeRF: Learn from One Look. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.

- [52] 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.
- [53] Chris Russell, Rui Yu, and Lourdes Agapito. Video Pop-up: Monocular 3D Reconstruction of Dynamic Scenes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014.
- [54] Yujiao Shi, Hongdong Li, and Xin Yu. Self-Supervised Visibility Learning for Novel View Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [55] Tomas Simon, Jack Valmadre, Iain Matthews, and Yaser Sheikh. Kronecker-Markov Prior for Dynamic 3D Reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, November 2017.
- [56] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. DeepVoxels: Learning Persistent 3D Feature Embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [57] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [58] Deqing Sun, Erik B. Sudderth, and Hanspeter Pfister. Layered RGBD scene flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [59] Zachary Teed and Jia Deng. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *Proceedings of the 16th European Conference on Computer Vision (ECCV)*, 2020.
- [60] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, August 2021.
- [61] Jack Valmadre and Simon Lucey. General trajectory prior for Non-Rigid reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2012.
- [62] Minh Vo, Srinivasa G. Narasimhan, and Yaser Sheikh. Spatiotemporal Bundle Adjustment for Dynamic 3D Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [63] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. IBRNet: Learning Multi-View Image-Based Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [64] 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 (CVPR)*, 2020.

- [65] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time Neural Irradiance Fields for Free-Viewpoint Video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [66] Wenpeng Xing and Jie Chen. Temporal-MPI: Enabling Multi-plane Images for Dynamic Scene Modelling via Temporal Basis Learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022.
- [67] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-NeRF: Point-based Neural Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.
- [68] Zexiang Xu, Sai Bi, Kalyan Sunkavalli, Sunil Hadap, Hao Su, and Ravi Ra. Deep View Synthesis from Sparse Photometric Images. *ACM transactions on graphics (TOG)*, 2019.
- [69] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. MVSNet: Depth Inference for Unstructured Multi-view Stereo. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [70] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [71] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural Radiance Fields From One or Few Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [72] Wentao Yuan, Zhaoyang Lv, Tanner Schmidt, and Steven Lovegrove. STaR: Self-supervised Tracking and Reconstruction of Rigid Objects in Motion with Neural Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [73] Xiaoshuai Zhang, Sai Bi, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. NeRFusion: Fusing Radiance Fields for Large-Scale Scene Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [74] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. HumanNeRF: Efficiently Generated Human Radiance Field from Sparse Inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.
- [75] Enliang Zheng, Dinghuang Ji, Enrique Dunn, and Jan-Michael Frahm. Sparse Dynamic 3D Reconstruction from Unsynchronized Videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, December 2015.
- [76] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. In *ACM SIGGRAPH*, August 2018.
- [77] C. Lawrence Zitnick, Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. High-quality video view interpolation using a layered representation. *ACM transactions on graphics (TOG)*, 2004.

- [78] Michael Zollhöfer, Matthias Nießner, Shahram Izadi, Christoph Rehmann, Christopher Zach, Matthew Fisher, Chenglei Wu, Andrew Fitzgibbon, Charles Loop, Christian Theobalt, and Marc Stamminger. Real-time non-rigid reconstruction using an RGB-D camera. *ACM transactions on graphics (TOG)*, July 2014.