

Gear-NeRF: Free-Viewpoint Rendering and Tracking with Motion-aware Spatio-Temporal Sampling

Liu, Xinhang; Tai, Yu-wing; Tang, Chi-Keung; Miraldo, Pedro; Lohit, Suhas; Chatterjee, Moitrya

TR2024-042 April 27, 2024

Abstract

Extensions of Neural Radiance Fields (NeRFs) to model dynamic scenes have enabled their near photo-realistic, free-viewpoint rendering. Although these methods have shown some potential in creating immersive experiences, two drawbacks limit their ubiquity: (i) a significant reduction in reconstruction quality when the computing budget is limited, and (ii) a lack of semantic understanding of the underlying scenes. To address these issues, we introduce Gear-NeRF, which leverages semantic information from powerful image segmentation models. Our approach presents a principled way for learning a spatio-temporal (4D) semantic embedding, based on which we introduce the concept of gears to allow for stratified modeling of dynamic regions of the scene based on the extent of their motion. Such differentiation allows us to adjust the spatiotemporal sampling resolution for each region in proportion to its motion scale, achieving more photo-realistic dynamic novel view synthesis. At the same time, almost for free, our approach enables free-viewpoint tracking of objects of interest – a functionality not yet achieved by existing NeRF-based methods. Empirical studies validate the effectiveness of our method, where we achieve state-of-the-art rendering and tracking performance on multiple challenging datasets. The project page is available at: <https://merl.com/research/highlights/gear-nerf>.

IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2024

Gear-NeRF: Free-Viewpoint Rendering and Tracking with Motion-aware Spatio-Temporal Sampling

Xinhang Liu^{1†*} Yu-Wing Tai² Chi-Keung Tang^{1†}
Pedro Miraldo³ Suhas Lohit³ Moitrey Chatterjee³

¹HKUST ²Dartmouth College ³Mitsubishi Electric Research Laboratories (MERL)

xliufe@connect.ust.hk, yu-wing.tai@dartmouth.edu, cktang@cse.ust.hk,
miraldo@merl.com, slohit@merl.com, metro.smiles@gmail.com

Abstract

Extensions of Neural Radiance Fields (NeRFs) to model dynamic scenes have enabled their near photo-realistic, free-viewpoint rendering. Although these methods have shown some potential in creating immersive experiences, two drawbacks limit their ubiquity: (i) a significant reduction in reconstruction quality when the computing budget is limited, and (ii) a lack of semantic understanding of the underlying scenes. To address these issues, we introduce **Gear-NeRF**, which leverages semantic information from powerful image segmentation models. Our approach presents a principled way for learning a spatio-temporal (4D) semantic embedding, based on which we introduce the concept of gears to allow for stratified modeling of dynamic regions of the scene based on the extent of their motion. Such differentiation allows us to adjust the spatio-temporal sampling resolution for each region in proportion to its motion scale, achieving more photo-realistic dynamic novel view synthesis. At the same time, almost for free, our approach enables free-viewpoint tracking of objects of interest – a functionality not yet achieved by existing NeRF-based methods. Empirical studies validate the effectiveness of our method, where we achieve state-of-the-art rendering and tracking performance on multi-view challenging datasets. The project page is available at: <https://merl.com/research/highlights/gear-nerf>.

1. Introduction

Reconstructing 3D scenes has a broad range of applications, including Virtual Reality/Augmented Reality (VR/AR), 3D animation, game production, and film creation which allow users to observe scenes from any desired viewpoint. While it is crucial to reconstruct *static* scenes, towards which significant progress has been made, it is even more crucial to

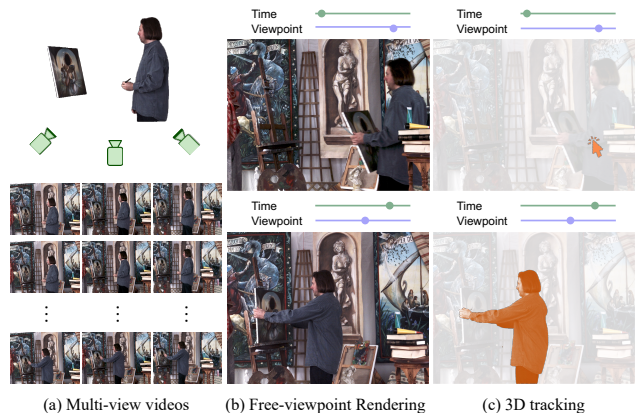


Figure 1. (a) Our method takes RGB videos captured from a camera array as input. (b) Trained Gear-NeRF achieves photo-realistic real-time free-viewpoint rendering of a dynamic scene. (c) With users giving a single click at any time and from any viewpoint, our method can perform free-viewpoint tracking of the target object.

reconstruct *dynamic* scenes, as the world around us is characterized by a constant state of flux, with many objects in it - in a state of motion.

Recent advances in novel view synthesis, such as Neural Radiance Fields (NeRFs) [57] have inspired numerous studies to extend them to dynamic 3D scenes. Existing approaches either employ a deformation field to map neural fields from a given time to a canonical space [22, 46, 61, 62, 66, 100], or directly model dynamic scenes as a 4D space-time grid [1, 6, 26]. Though these methods offer improved rendering quality by utilizing more accessible inputs compared to previous solutions [21, 43, 44, 78], they still struggle to ensure rendering quality in low-resource settings, requiring carefully engineered efforts. Further, most dynamic radiance field approaches adopt a naive spatio-temporal sampling strategy, without discerning the different scales of motion across different regions in the scene.

We propose to fix this issue by leveraging a semantic understanding of dynamic scenes. Intuitively, a reconstruction system aware of the distinction between static and dynamic regions in a scene can perform more focused

*Work mainly done when XL was an intern at MERL.

†XL and CT are supported in part by the Research Grant Council of the Hong Kong SAR grant no. 16201420.

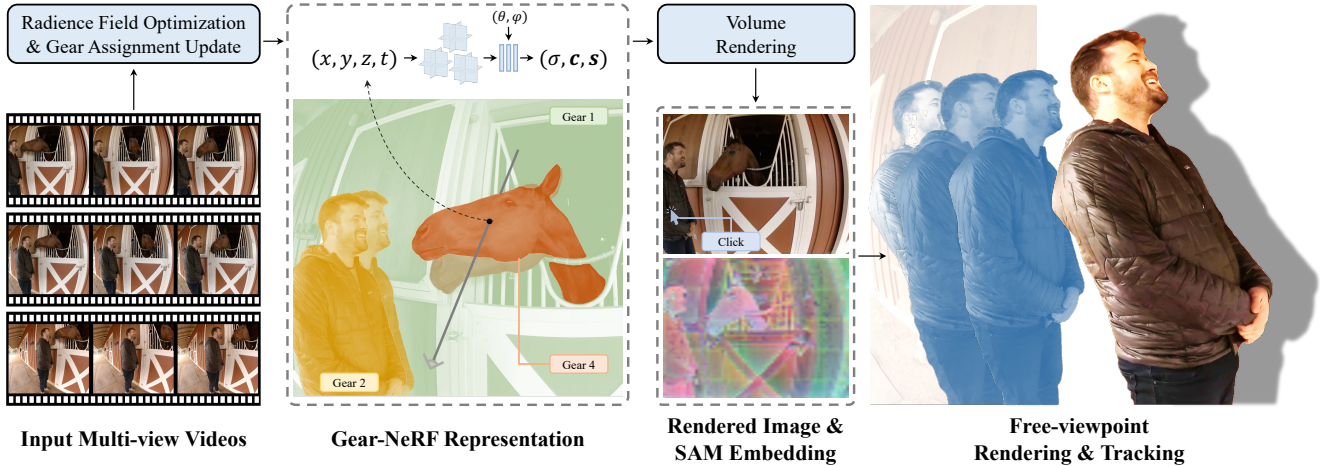


Figure 2. **Pipeline of Gear-NeRF:** Gear-NeRF takes multi-view videos as input. After optimizing the serial 4D feature volumes (Section 4.1), it maps space-time coordinates to a 4D semantic embedding (Section 4.2), in addition to the volume density and view-dependent radiance color. Regions with larger motion are automatically assigned higher gear levels (Section 4.3) and as a result, receive higher-resolution spatio-temporal sampling (Section 4.4). Furthermore, Gear-NeRF is capable of performing free-viewpoint tracking of a target object with prompts as simple as a user click (Section 4.5).

sampling in the dynamic regions, which inherently require more resources per unit volume than static regions, due to their time-evolving nature. Accordingly, dynamic regions can be further stratified according to their scale of motion. To this end, this paper presents *Gear-NeRF*, a framework that leverages semantic embedding from powerful image segmentation models [39] for stratified modeling of 4D scenes. Gear-NeRF optimizes for a 4D semantic embedding, based on which we introduce the concept of *gear* to smartly determine the appropriate region-specific resolution of spatio-temporal sampling in the NeRF. Regions with larger motion scales are assigned higher gears, through our gear determination scheme and we accordingly perform higher-resolution spatio-temporal sampling. Empirical studies reveal that this motion-aware sampling strategy improves the quality of synthesized images, over competing approaches. As a by-product of our semantically embedded representation, we achieve free-viewpoint object tracking, given user prompts. Figure 1 presents an overview of the capabilities of our method.

Gear-NeRF makes two primary advancements: (i) enhanced dynamic novel view synthesis by resorting to smarter spatio-temporal sampling, and (ii) the ability for free-viewpoint tracking of objects of interest. The latter is a capability not yet realized by existing NeRF methods for dynamic scenes. We perform extensive experiments on multiple datasets to validate the generalizability and robustness of our method, which shows state-of-the-art performances for both tasks across all datasets.

2. Related Work

Neural Radiance Fields: NeRF [57] is a recent breakthrough among novel view synthesis methods that uses multilayer perceptrons (MLPs) to parameterize the appearance

and density for each point in 3D space, given any viewing direction of the scene. Researchers have extended NeRF along various dimensions [79], including improving rendering quality [2–4, 12, 18, 31], handling challenging conditions such as large scenes [56, 69, 76], view-dependent appearances [30, 53, 81], and sparse inputs [34, 52, 60, 87, 88, 94]. NeRF-like neural representations have also found applications in semantic segmentation [41, 51, 103] and 3D content generation [9, 10, 50, 65]. Recent work has shown that replacing the deep MLPs with a feature voxel grid can significantly improve training and inference speed [11, 25, 59, 75]. On the other hand, a more recent approach to further improve visual quality, rendering time, and performance entails representing the scene with 3D Gaussians [36]. Our approach, while drawing upon many of these approaches, deals with dynamic scenes which is beyond the scope of these methods.

Neural Representations for Dynamic Scenes: NeRF-like representations have recently been extended to model dynamic scenes in high fidelity [17, 28, 33, 40, 47–49, 54, 64, 67, 85, 90, 91, 98]. One straightforward approach to do this is to directly condition the radiance field on time [27, 45, 46, 89]. Alternatively, several methods model a deformation field to map coordinates from different time stamps to a common canonical space [22, 24, 61, 62, 66, 80, 99, 100]. Some recent approaches [6, 26, 72, 82] represent the scene using a 4D space-time grid, which is decomposed into sets of planar representations for training efficiency. Other techniques for improving rendering fidelity and frame rate include Fourier PlenOctrees [83], ray-conditioned sample prediction networks [1], 4D space decomposition (static/dynamic/newly appeared regions) [74], and explicit voxel grids [24]. 3D Gaussians have also been adapted to model dynamic scenes [55, 86, 96, 97]. While these approaches paved the initial path for rendering dy-

dynamic scenes, semantically aware modeling of the scene is absent, a caveat that our proposed method seeks to address.

Segment Anything Model (SAM): SAM [39] is a powerful promptable image segmentation model, which showcases remarkable zero-shot generalization abilities and can produce semantically consistent masks, given a single foreground point on the image. HQ-SAM [35] is an improvement on SAM that enhances the quality of masks, especially on objects with intricate boundaries and structures. Recent works [20, 93] have extended SAM to perform interactive video object segmentation. These methods utilize SAM for mask initialization or correction and then employ state-of-the-art mask trackers [19, 95] for mask tracking and prediction [14]. Recent methods have also leveraged SAM for tracking multiple reference objects in a video [13, 102]. This work uses SAM to profile the scene into semantic regions, which are then grouped based on motion scales.

3D Semantic Understanding: Existing methods for 3D visual understanding [15, 32, 77, 92] mainly focus on closed set segmentation of point clouds or voxels. NeRF’s ability to integrate information across multiple views has led to its applications in 3D semantic segmentation [103], object segmentation [23, 51, 58], panoptic segmentation [73], and interactive segmentation [29, 70]. Kobayashi *et al.* [41] explored the effectiveness of embedding pixel-aligned features [7, 42] into NeRFs for 3D manipulations. LERF [37] fuses CLIP embeddings [68] and NeRFs to enable language-based localization in 3D NeRF scenes. Recent work has enabled click/text-based 3D segmentation by learning a 3D SAM embedding [16] or inverse rendering of SAM-generated masks [8]. We, on the other hand, seek to utilize the synergy of dynamic NeRFs and SAM segments to derive a semantic understanding of a dynamic 3D scene for tracking objects of interest in novel views – a first of its kind effort.

3. Preliminaries

Neural Radiance Fields (NeRFs): Vanilla NeRFs [57] employ a multi-layer perceptron (MLP) with sinusoidal positional encoding to map a 3D-spatial coordinate $\mathbf{x} = (x, y, z)$ and a viewing direction $\mathbf{d} = (\theta, \phi)$ to a volume density $\sigma \in [0, 1]$ and an emitted RGB, $\mathbf{c} \in \mathbb{R}^3$. Rendering each image pixel involves casting a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ from the camera center \mathbf{o} through the pixel along direction \mathbf{d} . The predicted color for the corresponding pixel is computed as:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i \mathbf{c}_i, \quad (1)$$

where $T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right)$, $\alpha_i = 1 - \exp(-\sigma_i \delta_i)$, and $\delta_j = t_{j+1} - t_j$. A vanilla NeRF is trained by minimizing the mean squared error between the input images and the predicted images, obtained by rendering the scene from the viewpoints from which the input images have been

captured, with the training loss given by:

$$\mathcal{L}_{\text{pho}} = \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}(\mathbf{r})\|_2^2. \quad (2)$$

where \mathcal{R} is the set of all rays projected from the input image.

Planar-Factorized 4D Volumes: A recent emerging trend of handling dynamics using radiance field representations is to directly adapt them to be conditioned on a frame index t (denoting time) in addition to \mathbf{x} and \mathbf{d} . This can be accomplished by learning a mapping from $(\mathbf{x}, \mathbf{d}, t)$ to (σ, \mathbf{c}) using planar-factorized 4D volumes [1, 6, 26, 72]. These methods attempt to learn a 4D feature vector for every (\mathbf{x}, t) , by projecting it to a set of 2D-planes. Embeddings of these projections on these planes can then be integrated to obtain the embedding for the 4D point. This can be mathematically represented as follows:

$$\begin{aligned} \mathbf{f}(\mathbf{x}, t) = & \mathbf{B}_1(\mathbf{h}_1(x, y) \odot \mathbf{k}_1(z, t)) \\ & + \mathbf{B}_2(\mathbf{h}_2(x, z) \odot \mathbf{k}_2(y, t)) \\ & + \mathbf{B}_3(\mathbf{h}_3(y, z) \odot \mathbf{k}_3(x, t)). \end{aligned} \quad (3)$$

where $\mathbf{h}_i(\cdot, \cdot)$ and $\mathbf{k}_i(\cdot, \cdot)$ are functions (evaluated by bilinear interpolation on regularly spaced 2D feature grids) embedding coordinate tuples to features of dimension M , “ \odot ” denotes an element-wise product, and $\mathbf{B}_i(\cdot)$ denotes a linear transform which maps the products to feature vectors. Subsequently, a tiny MLP can map the feature vector $\mathbf{f}(\cdot, \cdot)$ to the volume density, σ , and the view-dependent emitted color, \mathbf{c} , given the viewing direction \mathbf{d} .

4. Proposed Method

Given a set of W input videos, $\mathcal{V} = \{V_1, V_2, \dots, V_W\}$ of a dynamic scene, with calibrated camera poses, our approach represents the scene using a series of 4D feature volumes (Section 4.1) along with 4D semantic embeddings (Section 4.2).

Analogous to multiple gears in motor vehicles for optimizing engine performance, Gear-NeRF stratifies this semantically embedded scene representation into N_{gear} levels, based on the motion scales. Each of these levels is called a *gear*. Through our training scheme, regions with larger motion are assigned higher gear levels (Section 4.3) and as a result, are more densely sampled (Section 4.4) for improved dynamic novel view synthesis. Our 4D semantic embedding also enables a new functionality, almost for free – free-viewpoint tracking of target objects, given simple user prompts like clicks (Section 4.5). Figure 2 shows the overall pipeline of *Gear-NeRF*.

4.1. Serial 4D Feature Volumes

Instead of using a unified 4D volume to represent a dynamic scene [1, 6, 26, 72], our representation consists of a series of feature volumes, each corresponding to a gear level, \mathcal{G} . Specifically, for any space-time coordinate (\mathbf{x}, t) , its feature

vector corresponding to \mathcal{G} is computed as follows:

$$\begin{aligned} \mathbf{f}^{\mathcal{G}}(\mathbf{x}, t) &= \mathbf{B}_1(\mathbf{h}_1(x, y) \odot \mathbf{k}_1^{\mathcal{G}}(z, t)) \\ &+ \mathbf{B}_2(\mathbf{h}_2(x, z) \odot \mathbf{k}_2^{\mathcal{G}}(y, t)) \\ &+ \mathbf{B}_3(\mathbf{h}_3(y, z) \odot \mathbf{k}_3^{\mathcal{G}}(x, t)). \end{aligned} \quad (4)$$

The vector-valued functions $\mathbf{h}_j(\cdot, \cdot)$ and linear transforms $\mathbf{B}_j(\cdot)$ are shared by all gears, while each gear has its own spatio-temporal embedding $\mathbf{k}_j^{\mathcal{G}}(\cdot, \cdot)$, in M -dimensional space. Therefore, each gear describes regions of a certain scale of motion while the purely spatial features can be shared among all gears.

We obtain the gear level at any spatio-temporal coordinate also from a planar-factorized 4D feature volume. Specifically, the gear level at (\mathbf{x}, t) is computed as:

$$\begin{aligned} g(\mathbf{x}, t) &= \mathbf{1}^\top (\mathbf{h}'_1(x, y) \odot \mathbf{k}'_1(z, t)) \\ &+ \mathbf{1}^\top (\mathbf{h}'_2(x, z) \odot \mathbf{k}'_2(y, t)) \\ &+ \mathbf{1}^\top (\mathbf{h}'_3(y, z) \odot \mathbf{k}'_3(x, t)), \end{aligned} \quad (5)$$

where $\mathbf{1}$ is a vector of ones, $\mathbf{h}'_i(\cdot, \cdot)$ and $\mathbf{k}'_i(\cdot, \cdot)$ are M -dimensional embedding functions. This however defines a continuous feature volume. To map it to the gear level integers, we apply the following projection operation:

$$p(\mathbf{x}, t) = \begin{cases} 1, & \text{if } g(\mathbf{x}, t) < 1, \\ N_{\text{gear}}, & \text{if } g(\mathbf{x}, t) \geq N_{\text{gear}}, \\ \lceil g(\mathbf{x}, t) \rceil, & \text{otherwise.} \end{cases} \quad (6)$$

Based on this gear level volume, we define a 4D mask for a region at gear level \mathcal{G} as:

$$m_{\mathcal{G}}(\mathbf{x}, t) = \begin{cases} 1, & \text{if } p(\mathbf{x}, t) = \mathcal{G}, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

The final feature vector at (\mathbf{x}, t) is computed as:

$$\mathbf{f}(\mathbf{x}, t) = \sum_{\mathcal{G}=1}^{N_{\text{gear}}} m_{\mathcal{G}}(\mathbf{x}, t) \mathbf{f}^{\mathcal{G}}(\mathbf{x}, t). \quad (8)$$

Subsequently, a tiny MLP, F_θ , maps these feature vectors $\mathbf{f}(\cdot, \cdot)$ as well as the viewing direction \mathbf{d} to the volume density σ and radiance color \mathbf{c} . This allows us to obtain a photometric rendering of the scene.

4.2. 4D Semantic Embedding

Gear-NeRF leverages the strong object priors of the SAM [39] model to acquire a semantic understanding of the scene, for improved photometric rendering (Section 4.3 and Section 4.4) as well as free-viewpoint object tracking (Section 4.5). Toward this end, we utilize the SAM encoder to obtain 2D feature maps from the frames of each video. We then optimize a 4D SAM embedding field by supervising it with these 2D feature maps. In particular, the MLP above, F_θ , is configured to output a 4D semantic embedding \mathbf{s} for




Figure 3. **Illustration of Gear Assignment Update:** For each gear assignment update, we calculate the **rendering loss map** between the rendered RGB-SAM map and the ground truth and identify the centers of the patches with the maximum and minimum losses, marked in red and green (second column). These points are then fed into the SAM decoder as positive and negative prompts to generate an **upshift mask** representing the areas that need to be shifted to a higher gear (last column). After the first gear assignment update, we see that the next candidate region for upshift is situated where the horse is located, and so on. Upshift mask colors imply the gear levels after the update (blue-2, green-3, red-4).

a given space-time coordinate in addition to the density, σ , and color, \mathbf{c} . To render 2D semantic feature maps in a given view, we compute the semantic feature of a pixel in the feature map by tracing a ray through it and perform volume rendering analogous to Equation 1, as follows:

$$\hat{\mathbf{S}}(\mathbf{r}) = \sum_{i=1}^N T_i \alpha_i \mathbf{s}_i. \quad (9)$$

This SAM embedding is supervised by minimizing the mean squared error between the prediction and the ground truth features ($\mathbf{S}(\mathbf{r})$) from the SAM encoder, as shown:

$$\mathcal{L}_{\text{SAM}} = \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{\mathbf{S}}(\mathbf{r}) - \mathbf{S}(\mathbf{r})\|_2^2. \quad (10)$$

4.3. Training Scheme with Gear Assignment

With gear initialization $g(\mathbf{x}, t) = 1, \forall \mathbf{x}, t$, the (semantically embedded) radiance field optimization and gear assignment updating take place in an alternating fashion.

Gear Assignment Update: As illustrated in Figure 3, when updating the gear assignment after a period of radiance field

optimization, we find the regions rendered most poorly from the rendering loss maps and increment their gears for denser spatio-temporal sampling. The following steps lay out the process for updating gear assignments to regions:

- We sample a number of viewpoints and time steps and render 2D-images/SAM features for it. For every rendered RGB-SAM map, we compute a rendering loss map. Each pixel of the rendering loss map is computed as: $\mathcal{L}(\mathbf{r}) := \mathcal{L}_{\text{pho}}(\mathbf{r}) + \lambda \mathcal{L}_{\text{SAM}}(\mathbf{r})$. See Figure 3 for example loss maps.
- Next, we patchify each rendering loss map to find patches with the top- k largest/smallest average loss. The center coordinate of these patches serve as positive/negative prompts for the next step. See Figure 3 for example positive (red) / negative (green) prompts.
- We then feed the ground truth RGB images together with positive and negative prompts into the SAM decoder [39] to estimate an *upshift mask*. These masks tend to cover regions that have motions and are not satisfactorily rendered with the current sampling resolution. We have multiple upshift masks at different viewpoints and time steps.
- For every pixel of an upshift mask, we trace a ray and sample points along it and update the gear assignment by pushing $g(\mathbf{x}, t)$ towards incremented values.

In particular, in the last step, for each pixel within an upshift mask, a corresponding ray is traced that connects it with the camera center \mathbf{o} , along direction \mathbf{d} . Next, a set of points are sampled along this ray. The collection of sampled points, that lie on the rays emanating from within the masked region constitutes the set $\mathcal{S}_{\text{upshift}} = \{(\mathbf{x}_i^{\text{upshift}}, t_i^{\text{upshift}})\}_{i=1}^{N_{\text{upshift}}}$, where N_{upshift} represents the total count of points sampled from rays pertaining to the masked region. We then follow a similar procedure to sample a set of points pertaining to the unmasked region. We denote this set as: $\mathcal{S}_{\text{stay}} = \{(\mathbf{x}_i^{\text{stay}}, t_i^{\text{stay}})\}_{i=1}^{N_{\text{stay}}}$, with N_{stay} indicating the total number of points sampled from rays that pertain to the unmasked area. Next, for each sample point in each of $\mathcal{S}_{\text{upshift}}$ and $\mathcal{S}_{\text{stay}}$, we query its current gear level $p(\mathbf{x}, t)$. In order to assign new gear levels, we need to update the gear assignment function $g(\cdot, \cdot)$, which proceeds with the objective function:

$$\begin{aligned} \mathcal{L}_{\text{upshift}} = & \frac{1}{N_{\text{upshift}}} \sum_{(\mathbf{x}, t) \in \mathcal{S}_{\text{upshift}}} \|g(\mathbf{x}, t; \Theta) - (p(\mathbf{x}, t) + 1)\|_2^2 \\ & + \frac{\lambda_{\text{stay}}}{N_{\text{stay}}} \sum_{(\mathbf{x}, t) \in \mathcal{S}_{\text{stay}}} \|g(\mathbf{x}, t; \Theta) - p(\mathbf{x}, t)\|_2^2, \end{aligned} \quad (11)$$

where Θ denotes the set of optimizable parameters for $g(\cdot, \cdot)$. The minimization of the first term encourages sample points within the masked region $\mathcal{S}_{\text{upshift}}$ to have a gear level equal to their current gear incremented by one, resulting in an upshift of gear. Conversely, minimizing the second term encourages the points in $\mathcal{S}_{\text{stay}}$ to maintain their gear values, thereby encouraging the remaining regions to keep their current gear levels and avoiding unwanted up-

shifts. We update Θ via a single step of gradient descent as follows:

$$\Theta \leftarrow \Theta - \alpha \nabla_{\Theta} \mathcal{L}_{\text{upshift}}, \quad (12)$$

where α is the learning rate. Since, $g(\cdot, \cdot)$ is essentially derived from the embedding functions, $\mathbf{h}'_i(\cdot, \cdot)$ and $\mathbf{k}'_i(\cdot, \cdot)$ for $i \in \{1, 2, 3\}$, the aforementioned optimization step amounts to updating these embedding functions. Once updated, we proceed with a fresh round of gear assignment to increase the spatio-temporal sampling resolution for the regions that end up at a higher gear level than before.

Radiance Field Optimization: With the updated gear assignment, we increase the resolution of spatio-temporal sampling for the gear-shifted regions (Section 4.4) and then resume optimizing the radiance field. We alternate between the two processes: radiance field optimization (each time for L epochs), and gear assignment updates until the average variance of each rendering loss map is below a predetermined threshold. After this, we optimize the radiance field for an additional L' epochs without further gear assignment updates.

4.4. Motion-aware Spatio-Temporal Sampling

In this subsection, we explain our motion-aware spatio-temporal sampling strategy based on assigned gears, permitting differential processing of regions at different gear levels. By temporal sampling, we imply the choice of temporal resolution for planar-factorized 4D feature volumes, and by spatial sampling, we mean the strategy used to choose sampling points along each ray for volume rendering.

Motion-aware Temporal Sampling: To handle the increasing intensity of object motion, as reflected by their growing gear levels, we increment the temporal resolution for voxel grids. Specifically, $\mathbf{k}_j^{\mathcal{G}}$ in Equation 4 has increasing resolution along the time axis, thereby empowering the 4D feature volumes to better model the dynamics along the temporal axis. This ensures fast-moving objects can be more faithfully modeled without unsightly blurring. The temporal resolution for each gear’s feature volume is determined by linear interpolation between 1 (for $\mathcal{G} = 1$) and the total number of frames (for $\mathcal{G} = N_{\text{gear}}$).

Motion-aware Spatial Sampling: While denser sampling of points can improve reconstruction accuracy, increasing the number of sampling points throughout the scene can lead to prohibitive computational costs. Therefore, we propose a 3D point-splitting strategy as illustrated in Figure 4. We begin by sampling a relatively small number, n , of samples along each ray, assuming it is at the lowest gear level. If a sampled point belongs to a region with a higher gear, as determined by $p(\mathbf{x}, t)$, we then sample more densely in that region. For every sampled point in that region, we split it into $2^{p(\mathbf{x}, t)-1}$ points, equally spaced within the corresponding ray segment (at that gear level).




Figure 4. **Motion-aware Spatial Sampling:** We split each sampled point into $2^{p(x,t)}$ points, equally spaced within the corresponding ray segment. The top row shows the vanilla uniformly sampled points, while the bottom one shows the densely sampled points after splitting using our strategy.

4.5. Free-Viewpoint Tracking with User Prompts

Our 4D SAM embedding enables another useful functionality, almost for free – free-viewpoint object tracking, where the user only needs to provide as few as one click to extract the target object based on the 4D embedding. Next, we describe how, given a user-supplied point click at any arbitrary viewpoint and time step, we obtain the corresponding object mask at a novel viewpoint and time step.

Masks for Novel Viewpoints: The first step for this task entails finding the 3D correspondence of the user click. We trace a ray through the selected pixel, and by utilizing the volume density, we determine the depth at which the ray intersects with the first object surface it encounters. This yields the 3D coordinates of the point of intersection. Subsequently, the 3D coordinates of this intersection can be easily mapped into a 2D coordinate within any novel viewpoint image, using the camera pose of the new viewpoint. Alongside the rendered SAM feature map of the novel view, we feed this coordinate into the SAM decoder to generate the object mask for the novel view.

Masks for Novel Time Steps: For this task, we propagate an object mask to its neighboring time step. Specifically, with an object mask for a specific frame t , we calculate the bounding box of this mask and use this bounding box as a prompt to SAM for neighboring frames $t' = t + 1$ or $t' = t - 1$. By inputting this prompt along with the rendered SAM feature map at t' into the SAM decoder, we can obtain the object mask for t' . Combining the above two processes, we can start from a single click and get the object mask in any viewpoint and time step.

5. Experiments

We assess the performance of our proposed Gear-NeRF for dynamic novel view synthesis and free-viewpoint object tracking across a range of challenging datasets, comparing it with state-of-the-art methods. Through ablation studies, we provide empirical evidence of the effectiveness of its fundamental components. We kindly refer the reader to our supplementary material for additional experimental details and results, including videos for free-viewpoint rendering and tracking.

5.1. Experimental Setup

Implementation Details: We implement our method using PyTorch [63] and conduct experiments on an NVIDIA RTX 4090 GPU with 24 GB RAM. We divide each input video into chunks of 100 frames. For every chunk, we train a model for approximately 2.5 hours. Our 4D feature volumes yield embeddings with a dimension of $M = 32$. We set the gear number N_{gear} to 4. We find patches with top- $k = 3$ largest/smallest average loss for gear assignment updates to obtain prompts. The rendering loss map is computed with $\lambda = 0.01$. For the optimization of radiance fields, $L = 3$ and $L' = 10$. In our motion-aware spatial sampling, each ray initially has $n = 64$ sampling points. We use an initial learning rate of 0.02 for all the parameters and optimize them using ADAM [38]. For Equation 12, we use $\alpha = 0.02$.

Datasets: (i) **The Technicolor light field dataset** [71] includes diverse indoor environment videos captured by a 4×4 camera rig. We evaluate on 4 sequences (*Train*, *Theater*, *Painter*, *Birthday*) at the original 2048×1088 resolution, holding out the same view as prior work [1] (the second row and second column) for evaluation. (ii) **The Neural 3D Video dataset** [45] includes indoor multi-view video sequences captured by 20 cameras at a resolution of 2704×2028 pixels. We experiment on 6 sequences (*Cut Roasted Beef*, *Flame Steak*, *Coffee Martini*, *Cook Spinach*, *Flame Salmon*, *Sear Steak*), downsampling by a factor of 2 and holding out the central view (akin to prior work [1]) for evaluation. (iii) **The Google Immersive dataset** [5] contains light field videos of indoor and outdoor scenes captured by a time-synchronized 46-fisheye camera rig, with a resolution of 2560×1920 pixels. We experiment with 9 sequences from it (*Flames*, *Truck*, *Horse*, *Car*, *Welder*, *Exhibit*, *Face Paint 1*, *Face Paint 2*, *Cave*). We downsample the video resolution by a factor of 2 and hold out the central view (like prior work [1]) for evaluation. For our experiments, we adopted the same resolution and held-out view selection as prior work [1].

Evaluation Metrics: For the task of novel view synthesis of dynamic scenes, we evaluate using the following standard metrics: (i) Peak Signal-to-Noise Ratio (PSNR), (ii) Structural Similarity Index Measure (SSIM) [84] and (iii) Learned Perceptual Image Patch Similarity (LPIPS) [101], by comparing the reconstructed frames against ground truth images. These metrics are computed on the held-out view and averaged across all frames. For the free-viewpoint object tracking task, we designate a specific viewpoint and time step for the user to give prompting clicks. Subsequently, we assess the quality of the predicted object masks at novel viewpoints. The quality of the object mask is quantified in terms of the Mean Intersection over Union (mIoU) and accuracy (Acc.), which are calculated against the ground truth mask, manually annotated utilizing Adobe Photoshop. Additionally, we present the same metrics computed for *novel time steps*, denoted as t-mIoU and t-Acc.

Baselines: We run a comprehensive comparison of our




Figure 5. **Qualitative comparisons for novel view synthesis on the Technicolor dataset [71]:** We qualitatively compare our approach against HyperReel [1] and Neural 3D Video [45]. Our approach better recovers fine details like patterns on the toys or stripes on the shirt.

method against a range of recent NeRF-based baseline methods: (i) ST-NeRF [100], (ii) HexPlane [6], (iii) HyperReel [1], and (iv) MixVoxels [82]. As the first method to enable promptable free-viewpoint object tracking under the NeRF setting, there is no established direct baseline for this specific task. However, SA3D [8], a recent method for segmenting static scenes, is treated as a baseline for predicting masks at novel viewpoints corresponding to the *prompted time step*.

5.2. Results

Dynamic Novel View Synthesis: As shown in Figure 5, Gear-NeRF produces high-quality novel view synthesis of dynamic scenes, accurately modeling real-world scenes with intricate motions and fine details. For example, patterns on the toys or stripes on the shirt, are faithfully rendered, resulting in more photo-realistic images compared to all of the baselines. Table 1, presents quantitative comparisons of our method against the baselines. While Gear-NeRF has longer training (Tr. Time) / inference times (FPS) compared to some baselines, it almost always achieves the best performance in terms of rendering quality.

Free-Viewpoint Tracking: In Figure 6, we present the object masks obtained by our method based on the user prompts provided at a specified viewpoint and time step. Specifically, the first row displays masked objects at the prompted viewpoint and time step. The second row shows novel view masks at the prompted time step. The third row shows novel view masks at novel time steps. We see that masks obtained from Gear-NeRF show precise boundaries,

Table 1. **Quantitative comparisons for dynamic novel view synthesis:** Our method outperforms all baselines across all datasets on all metrics. We report means over all scenes for each dataset. **Best** and **second best** results are highlighted.

Dataset	Method	PSNR (↑)	SSIM (↑)	LPIPS (↓)	Tr. Time(↓)	FPS(↑)
Technicolor [71]	ST-NeRF [100]	30.86	0.883	0.101	344 min	0.7
	HyperReel [1]	31.04	0.887	0.092	97 min	7.7
	MixVoxels [82]	28.99	0.842	0.103	14 min	18.9
	Ours	32.21	0.919	0.058	148 min	7.9
Google Immersive [5]	HexPlane [6]	27.67	0.808	0.188	527 min	0.9
	HyperReel [1]	28.32	0.862	0.145	117 min	7.4
	MixVoxels [82]	27.14	0.835	0.209	19 min	16.8
	Ours	28.74	0.876	0.122	189 min	7.0
Neural 3D Video [45]	ST-NeRF [100]	31.03	0.890	0.081	367 min	0.7
	HyperReel [1]	31.12	0.928	0.065	129 min	6.1
	MixVoxels [82]	30.69	0.944	0.139	15 min	16.3
	Ours	31.80	0.936	0.058	204 min	6.8

compared to SA3D. Table 2 presents quantitative assessment of the quality of the masks generated of our method versus SA3D. Our method exceeds 90% across all metrics and datasets, demonstrating the effectiveness of our approach. Our gains over SA3D can be attributed to the fact that SA3D, as a static scene segmentation method, does not utilize information across all time frames, whereas our approach does. SA3D is incapable of predicting masks for novel time steps, and as such, the corresponding entries are marked as not applicable, a shortcoming which our method does not have.

5.3. Evaluations

To verify the effectiveness of the design choices of Gear-NeRF, we perform extensive ablation studies on the *Truck* scene of the Google Immersive dataset [5].

Motion-aware Temporal Sampling: An intuitive strategy for adjusting the temporal sampling is to directly modify



Figure 6. **Qualitative comparisons of free-viewpoint object tracking on Technicolor [71] and Neural 3D Video [45] datasets:** Our method can obtain desirable object masks, with clear edges, from prompting points provided by users at desired time steps and viewpoints.

Table 2. **Quantitative comparisons for free-viewpoint tracking:** t-mIoU and t-Acc are metrics used for evaluating novel view masks at novel time steps, not applicable to SA3D. Reported metrics are averages over all scenes for each dataset.

Dataset	Method	mIoU (\uparrow)	Acc. (\uparrow)	t-mIoU (\uparrow)	t-Acc. (\uparrow)
Technicolor [71]	SA3D [8]	96.4	97.1	N/A	N/A
	Ours	97.4	97.6	92.1	93.3
Google Immersive [5]	SA3D [8]	94.1	94.8	N/A	N/A
	Ours	94.3	95.0	91.5	92.8
Neural 3D Video [45]	SA3D [8]	93.1	94.0	N/A	N/A
	Ours	93.4	94.3	90.6	92.3

the temporal resolution of the 4D feature volume and make it uniform across all regions of the scene. In Table 3, we demonstrate the performance of this naive temporal sampling strategy using medium (25) or dense (100) temporal resolutions on 100-frame input videos. The results show that increasing the temporal resolution universally does not necessarily yield better performance. In contrast, our proposed method, which utilizes motion-aware temporal sampling, achieves the best results. We surmise that this may be attributed to the model distributing the volume’s capacity sparsely across a large number of time steps, which is not the case for our method.

Motion-aware Spatial Sampling: To validate the effectiveness of our motion-aware spatial sampling strategy, we compare it against other sampling strategies. The first variant involves uniform sampling along the ray, and the second uses a sample prediction network (SPN) like Attal *et al.* [1] for sampling. These variants sample 128 points on each ray. Results in Table 3 show that our full model employing motion-aware spatial sampling outperforms these variants, perhaps by having a better sense of where to sample more from, derived from its semantic-aware embedding.

SAM embedding: To verify that introducing SAM embedding can improve rendering quality (rather than just enabling segmentation or tracking), we tested a variant of Gear-NeRF without SAM embedding (it instead thresholds loss maps of rendered RGB frames to obtain upshift masks). As shown in Table 3, the absence of SAM em-

Table 3. **Ablation Study:** Ablations on our spatio-temporal sampling strategy and the number of gears (Truck/Google Immersive). **Best** and **second best** results are highlighted.

Method	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)
Naive temporal sampling (medium)	26.93	0.778	0.144
Naive temporal sampling (dense)	26.85	0.760	0.161
Naive spatial sampling	24.80	0.734	0.222
SPN spatial sampling	26.54	0.787	0.162
Ours (w/o embedding)	27.10	0.831	0.177
Ours ($N_{\text{gear}} = 2$)	26.93	0.785	0.166
Ours ($N_{\text{gear}} = 3$)	27.14	0.875	0.145
Ours ($N_{\text{gear}} = 4$)	27.49	0.892	0.136
Ours ($N_{\text{gear}} = 5$)	<u>27.46</u>	0.901	0.131

bedding for guiding gear assignment reduces the model’s rendering quality.

Number of Gears: We ablate on the numbers of gears. Increasing the number of gears allows for more fine-grained motion-aware spatio-temporal sampling, while increasing the computational cost. As shown in Table 3, a choice of up to 4 gear levels seems optimal, further increasing the number of gears does not result in significant improvements.

6. Conclusions

In this work, we introduced *Gear-NeRF*, an extension of dynamic NeRFs that leverages semantic information from powerful segmentation models for stratified modeling of dynamic scenes. Our approach learns a 4D (spatio-temporal) semantic embedding and introduces the concept of “gears” for differentiated modeling of scene regions based on their motion intensity. With determined gear assignments, Gear-NeRF adaptively adjusts its spatio-temporal sampling resolution to improve the photo-realism of rendered views. At the same time, Gear-NeRF provides the new functionality of free-viewpoint object tracking with user prompts as simple as a click. Our empirical studies underscore the effectiveness of Gear-NeRF, showcasing state-of-the-art performance in both rendering quality and object tracking across multiple challenging datasets.

References

- [1] Benjamin Attal, Jia-Bin Huang, Christian Richardt, Michael Zollhoefer, Johannes Kopf, Matthew O’Toole, and Changil Kim. HyperReel: High-fidelity 6-DoF video with ray-conditioned sampling. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [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 *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5855–5864, 2021. [2](#)
- [3] 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 5470–5479, 2022.
- [4] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19697–19705, 2023. [2](#)
- [5] Michael Broxton, John Flynn, Ryan Overbeck, Daniel Erickson, Peter Hedman, Matthew Duvall, Jason Dourgarian, Jay Busch, Matt Whalen, and Paul Debevec. Immersive light field video with a layered mesh representation. *ACM Transactions on Graphics (TOG)*, 39(4):86–1, 2020. [6](#), [7](#), [8](#)
- [6] 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. [1](#), [2](#), [3](#), [7](#)
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers, 2021. [3](#)
- [8] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Chen Yang, Wei Shen, Lingxi Xie, Dongsheng Jiang, Xiaopeng Zhang, and Qi Tian. Segment anything in 3d with nerfs, 2023. [3](#), [7](#), [8](#)
- [9] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5799–5809, 2021. [2](#)
- [10] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16123–16133, 2022. [2](#)
- [11] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, pages 333–350. Springer, 2022. [2](#)
- [12] Anpei Chen, Zexiang Xu, Xinyue Wei, Siyu Tang, Hao Su, and Andreas Geiger. Factor fields: A unified framework for neural fields and beyond. *arXiv preprint arXiv:2302.01226*, 2023. [2](#)
- [13] Jiaben Chen and Huaizu Jiang. SportssloMo: A new benchmark and baselines for human-centric video frame interpolation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [3](#)
- [14] Jiaben Chen, Renrui Zhang, Dongze Lian, Jiaqi Yang, Ziyao Zeng, and Jianbo Shi. iquery: Instruments as queries for audio-visual sound separation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14675–14686, 2023. [3](#)
- [15] Xiaokang Chen, Kwan-Yee Lin, Chen Qian, Gang Zeng, and Hongsheng Li. 3d sketch-aware semantic scene completion via semi-supervised structure prior. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [3](#)
- [16] Xiaokang Chen, Jiaxiang Tang, Diwen Wan, Jingbo Wang, and Gang Zeng. Interactive segment anything nerf with feature imitation. *arXiv preprint arXiv:2305.16233*, 2023. [3](#)
- [17] Yue Chen, Xuan Wang, Xingyu Chen, Qi Zhang, Xiaoyu Li, Yu Guo, Jue Wang, and Fei Wang. Uv volumes for real-time rendering of editable free-view human performance. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16621–16631, 2023. [2](#)
- [18] Zhang Chen, Zhong Li, Liangchen Song, Lele Chen, Jingyi Yu, Junsong Yuan, and Yi Xu. Neurf: A neural fields representation with adaptive radial basis functions. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4182–4194, 2023. [2](#)
- [19] Ho Kei Cheng and Alexander G. Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model, 2022. [3](#)
- [20] Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. Segment and track anything. *arXiv preprint arXiv:2305.06558*, 2023. [3](#)
- [21] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (TOG)*, 34(4):1–13, 2015. [1](#)
- [22] 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. IEEE Computer Society, 2021. [1](#), [2](#)
- [23] Zhiwen Fan, Peihao Wang, Xinyu Gong, Yifan Jiang, Dejia Xu, and Zhangyang Wang. Nerf-sos: Any-view self-supervised object segmentation from complex real-world scenes. *International Conference on Learning Representations (ICLR)*, pages arXiv–2209, 2023. [3](#)
- [24] 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 *ACM Transactions on Graphics (SIGGRAPH Asia)*, pages 1–9, 2022. [2](#)
- [25] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5501–5510, 2022. [2](#)
- [26] 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. 1, 2, 3
- [27] 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 5712–5721, 2021. 2
- [28] Chen Geng, Sida Peng, Zhen Xu, Hujun Bao, and Xiaowei Zhou. Learning neural volumetric representations of dynamic humans in minutes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8759–8770, 2023. 2
- [29] Rahul Goel, Dhawal Sirikonda, Saurabh Saini, and P.J. Narayanan. Interactive Segmentation of Radiance Fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [30] Yuan-Chen Guo, Di Kang, Linchao Bao, Yu He, and Song-Hai Zhang. Nerfren: Neural radiance fields with reflections. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18409–18418, 2022. 2
- [31] Wenbo Hu, Yuling Wang, Lin Ma, Bangbang Yang, Lin Gao, Xiao Liu, and Yuewen Ma. Tri-miprf: Tri-mip representation for efficient anti-aliasing neural radiance fields. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19774–19783, 2023. 2
- [32] Qiangui Huang, Weiyue Wang, and Ulrich Neumann. Recurrent slice networks for 3d segmentation on point clouds. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [33] Mustafa Işık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Nießner. Humanrf: High-fidelity neural radiance fields for humans in motion. *ACM Transactions on Graphics (SIGGRAPH)*, 2023. 2
- [34] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5885–5894, 2021. 2
- [35] Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan Liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Segment anything in high quality. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 3
- [36] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)*, 42(4):1–14, 2023. 2
- [37] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lurf: Language embedded radiance fields. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 3
- [38] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 6
- [39] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2, 3, 4, 5
- [40] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Transactions on Graphics (SIGGRAPH)*, 2023. 2
- [41] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:23311–23330, 2022. 2, 3
- [42] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *International Conference on Learning Representations (ICLR)*, 2022. 3
- [43] Hao Li, Bart Adams, Leonidas J Guibas, and Mark Pauly. Robust single-view geometry and motion reconstruction. *ACM Transactions on Graphics (TOG)*, 28(5):1–10, 2009. 1
- [44] Hao Li, Linjie Luo, Daniel Vlasic, Pieter Peers, Jovan Popović, Mark Pauly, and Szymon Rusinkiewicz. Temporally coherent completion of dynamic shapes. *ACM Transactions on Graphics (TOG)*, 31(1):1–11, 2012. 1
- [45] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5521–5531, 2022. 2, 6, 7, 8
- [46] 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*, pages 6498–6508, 2021. 1, 2
- [47] 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. 2
- [48] Haotong Lin, Sida Peng, Zhen Xu, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Efficient neural radiance fields for interactive free-viewpoint video. In *ACM Transactions on Graphics (SIGGRAPH Asia)*, pages 1–9, 2022.
- [49] Haotong Lin, Sida Peng, Zhen Xu, Tao Xie, Xingyi He, Hujun Bao, and Xiaowei Zhou. Im4d: High-fidelity and real-time novel view synthesis for dynamic scenes. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 2023. 2
- [50] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9298–9309, 2023. 2
- [51] Xinhang Liu, Jiaben Chen, Huai Yu, Yu-Wing Tai, and Chi-Keung Tang. Unsupervised multi-view object segmentation using radiance field propagation. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:17730–17743, 2022. 2, 3
- [52] Xinhang Liu, Shiu-hong Kao, Jiaben Chen, Yu-Wing Tai, and Chi-Keung Tang. Deceptive-nerf: Enhancing nerf reconstruction using pseudo-observations from diffusion models. *arXiv preprint arXiv:2305.15171*, 2023. 2
- [53] Xinhang Liu, Yu-Wing Tai, and Chi-Keung Tang. Cleanerf: Reformulating nerf to account for view-dependent observations. *arXiv preprint arXiv:2303.14707*, 2023. 2
- [54] Yu-Lun Liu, Chen Gao, Andreas Meuleman, Hung-Yu Tseng, Ayush Saraf, Changil Kim, Yung-Yu Chuang, Johannes Kopf, and Jia-Bin Huang. Robust dynamic radiance

- fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13–23, 2023. 2
- [55] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023. 2
- [56] Andreas Meuleman, Yu-Lun Liu, Chen Gao, Jia-Bin Huang, Changil Kim, Min H Kim, and Johannes Kopf. Progressively optimized local radiance fields for robust view synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16539–16548, 2023. 2
- [57] 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)*, 2020. 1, 2, 3
- [58] Ashkan Mirzaei, Tristan Aumentado-Armstrong, Konstantinos G Derpanis, Jonathan Kelly, Marcus A Brubaker, Igor Gilitschenski, and Alex Levinshtein. Spin-nerf: Multiview segmentation and perceptual inpainting with neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20669–20679, 2023. 3
- [59] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 2
- [60] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5480–5490, 2022. 2
- [61] 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 5865–5874, 2021. 1, 2
- [62] 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)*, 2021. 1, 2
- [63] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019. 6
- [64] Sida Peng, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Representing volumetric videos as dynamic mlp maps. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4252–4262, 2023. 2
- [65] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *International Conference on Learning Representations (ICLR)*, 2023. 2
- [66] 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 10318–10327, 2021. 1, 2
- [67] Yi-Ling Qiao, Alexander Gao, Yiran Xu, Yue Feng, Jia-Bin Huang, and Ming C Lin. Dynamic mesh-aware radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 385–396, 2023. 2
- [68] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of Machine Learning Research (PMLR)*, pages 8748–8763, 2021. 3
- [69] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12932–12942, 2022. 2
- [70] Zhongzheng Ren, Aseem Agarwala, Bryan Russell, Alexander G. Schwing, and Oliver Wang. Neural volumetric object selection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [71] Neus Sabater, Guillaume Boisson, Benoît Vandame, Paul Kerbirou, Frederic Babon, Matthieu Hog, Remy Gendrot, Tristan Langlois, Olivier Bureller, Arno Schubert, et al. Dataset and pipeline for multi-view light-field video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 30–40, 2017. 6, 7, 8
- [72] 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, 3
- [73] Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Norman Müller, Matthias Nießner, Angela Dai, and Peter Kotschieder. Panoptic lifting for 3d scene understanding with neural fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9043–9052, 2023. 3
- [74] 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 (TVCG)*, 29(5): 2732–2742, 2023. 2
- [75] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5459–5469, 2022. 2
- [76] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8248–8258, 2022. 2

- [77] Jiayang Tang, Xiaokang Chen, Jingbo Wang, and Gang Zeng. Point scene understanding via disentangled instance mesh reconstruction. *European Conference on Computer Vision (ECCV)*, 2022. 3
- [78] Jonathan Taylor, Jamie Shotton, Toby Sharp, and Andrew Fitzgibbon. The vitruvian manifold: Inferring dense correspondences for one-shot human pose estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 103–110. IEEE, 2012. 1
- [79] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Wang Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. In *Computer Graphics Forum*, pages 703–735. Wiley Online Library, 2022. 2
- [80] 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 *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12959–12970, 2021. 2
- [81] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5490. IEEE, 2022. 2
- [82] Feng Wang, Sinan Tan, Xinghang Li, Zeyue Tian, and Huaping Liu. Mixed neural voxels for fast multi-view video synthesis. *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2, 7
- [83] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plenotrees for dynamic radiance field rendering in real-time. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13524–13534, 2022. 2
- [84] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing (TIP)*, 13(4):600–612, 2004. 6
- [85] Ziyang Wang, Giljoo Nam, Tuur Stuyck, Stephen Lombardi, Chen Cao, Jason Saragih, Michael Zollhöfer, Jessica Hodgins, and Christoph Lassner. Neuwigs: A neural dynamic model for volumetric hair capture and animation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8641–8651, 2023. 2
- [86] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xingang Wang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. 2
- [87] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P Srinivasan, Dor Verbin, Jonathan T Barron, Ben Poole, et al. Reconfusion: 3d reconstruction with diffusion priors. *arXiv preprint arXiv:2312.02981*, 2023. 2
- [88] Jamie Wynn and Daniyar Turmukhambetov. Diffusionerf: Regularizing neural radiance fields with denoising diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4180–4189, 2023. 2
- [89] 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 9421–9431, 2021. 2
- [90] Zhen Xu, Sida Peng, Haotong Lin, Guangzhao He, Jiaming Sun, Yujun Shen, Hujun Bao, and Xiaowei Zhou. 4k4d: Real-time 4d view synthesis at 4k resolution. *arXiv preprint arXiv:2310.11448*, 2023. 2
- [91] Zhiwen Yan, Chen Li, and Gim Hee Lee. Nerf-ds: Neural radiance fields for dynamic specular objects. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8285–8295, 2023. 2
- [92] Bo Yang, Jianan Wang, Ronald Clark, Qingyong Hu, Sen Wang, Andrew Markham, and Niki Trigoni. Learning object bounding boxes for 3d instance segmentation on point clouds, 2019. 3
- [93] Jinyu Yang, Mingqi Gao, Zhe Li, Shang Gao, Fangjing Wang, and Feng Zheng. Track anything: Segment anything meets videos, 2023. 3
- [94] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8254–8263, 2023. 2
- [95] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation, 2022. 3
- [96] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101*, 2023. 2
- [97] Zeyu Yang, Hongye Yang, Zijie Pan, Xiatian Zhu, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. *arXiv preprint arXiv:2310.10642*, 2023. 2
- [98] Zhengming Yu, Wei Cheng, Xian Liu, Wayne Wu, and Kwan-Yee Lin. Monohuman: Animatable human neural field from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16943–16953, 2023. 2
- [99] Wentao Yuan, Zhaoyang Lv, Tanner Schmidt, and Steven Lovegrove. Star: Self-supervised tracking and reconstruction of rigid objects in motion with neural rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13144–13152, 2021. 2
- [100] Jiakai Zhang, Xinhang Liu, Xinyi Ye, Fuqiang Zhao, Yanshun Zhang, Minye Wu, Yingliang Zhang, Lan Xu, and Jingyi Yu. Editable free-viewpoint video using a layered neural representation. *ACM Transactions on Graphics (TOG)*, 40(4):1–18, 2021. 1, 2, 7
- [101] 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
- [102] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*, 2023. 3

- [103] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15838–15847, 2021. [2](#), [3](#)