

# Feature4X: Bridging Any Monocular Video to 4D Agentic AI with Versatile Gaussian Feature Fields

Shijie Zhou<sup>1\*</sup> Hui Ren<sup>2\*</sup> Yijia Weng<sup>3</sup> Shuwang Zhang<sup>1</sup> Zhen Wang<sup>1</sup> Dejia Xu<sup>4</sup>  
 Zhiwen Fan<sup>4</sup> Suyu You<sup>5</sup> Zhangyang Wang<sup>4</sup> Leonidas Guibas<sup>3</sup> Achuta Kadambi<sup>1</sup>  
<sup>1</sup>UCLA <sup>2</sup>MIT <sup>3</sup>Stanford <sup>4</sup>UT Austin <sup>5</sup>DEVCOM ARL  
<https://feature4x.github.io/>

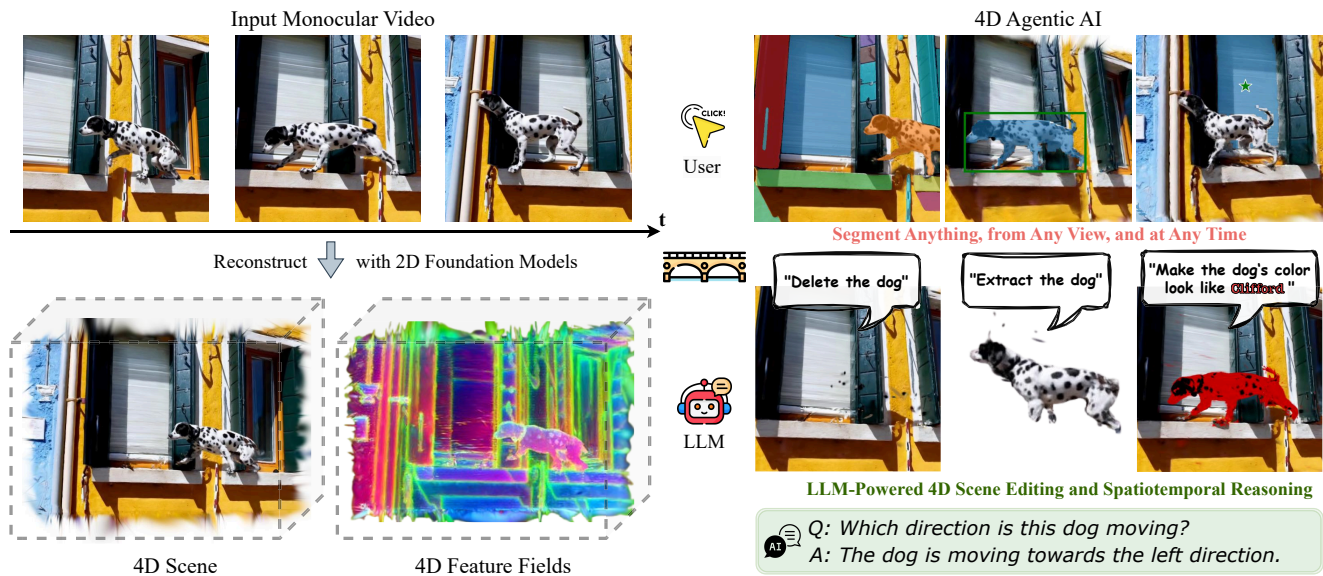


Figure 1. **Feature4X: Building 4D Interactive Scenes with Agentic AI from Monocular Videos.** By dynamically distilling model-conditioned features and integrating 2D foundation models with LLMs in feedback loops, Feature4X enables multimodal tasks across 2D, 3D, and 4D with high-level language inputs or direct user interactions, including (but not limited to) segmentation, scene editing, and VQA across novel views and all time steps, unlocking new possibilities for 4D agentic AI.

## Abstract

Recent advancements in 2D and multimodal models have achieved remarkable success by leveraging large-scale training on extensive datasets. However, extending these achievements to enable free-form interactions and high-level semantic operations with complex 3D/4D scenes remains challenging. This difficulty stems from the limited availability of large-scale, annotated 3D/4D or multi-view datasets, which are crucial for generalizable vision and language tasks such as open-vocabulary and prompt-based segmentation, language-guided editing, and visual question answering (VQA). In this paper, we introduce Feature4X, a universal framework designed to extend any functionality from 2D vision foundation model into the 4D realm, using only monocular video input, which is widely available from

user-generated content. The “X” in Feature4X represents its versatility, enabling any task through adaptable, model-conditioned 4D feature field distillation. At the core of our framework is a dynamic optimization strategy that unifies multiple model capabilities into a single representation. Additionally, to the best of our knowledge, Feature4X is the first method to distill and lift the features of video foundation models (e.g. SAM2, InternVideo2) into an explicit 4D feature field using Gaussian Splatting. Our experiments showcase novel view segment anything, geometric and appearance scene editing, and free-form VQA across all time steps, empowered by LLMs in feedback loops. These advancements broaden the scope of agentic AI applications by providing a foundation for scalable, contextually and spatiotemporally aware systems capable of immersive dynamic 4D scene interaction.

\*Equal contribution.

## 1. Introduction

The rapid evolution of 2D vision and multimodal models has been fueled by access to massive curated datasets with rich annotations, alongside breakthroughs spanning multiple domains. These developments have led to remarkable progress in various tasks, including image segmentation [64], image editing [6], promptable semantic segmentation [29, 66], and Visual Question Answering (VQA) [37, 81]. Despite these advances, the processing and interpretation of dynamic 3D data—crucial for applications such as autonomous driving, robotics, and 3D asset creation—still lag significantly behind the development of versatile and robust 2D foundation models. A fundamental challenge in 3D vision lies in processing multi-view images and widely available monocular videos, which often suffer from limited camera information and a scarcity of well-annotated, fine-grained per-frame datasets. To bridge this gap, we propose to leverage the latent features of well-trained representations from 2D and multimodal domains [29, 37, 64, 66, 81] and adaptively lift them into higher-dimensional feature fields (such as 3D and 4D) in a scalable and efficient manner, enabling the direct transfer of 2D functionalities into 4D with minimal data annotation and training overhead.

Constructing 4D feature fields from casually captured in-the-wild videos is significantly challenging. Existing approaches for static 3D feature fields [14, 28, 30, 33, 45, 63, 70, 78, 93, 98, 100] are not directly applicable to 4D scenarios for three primary reasons. First, these methods typically rely on well-calibrated, multi-view input images with precise camera poses, which are difficult to obtain from casual videos. Second, extending static feature fields to include an additional temporal dimension results in substantial memory demands, leading to unstable and costly optimization. Third, previous approaches are limited to task-specific feature fields, necessitating full-parameter re-training for adaptation to new tasks.

In this paper, we set an ambitious goal: to parse in-the-wild monocular videos into a unified dynamic 4D representation that not only reconstructs accurate appearance and geometry but also seamlessly integrates advanced functionalities from a range of 2D foundation models. To tackle the aforementioned challenges, we propose enhancing the dynamic 3D Gaussian Splatting-based 4D scene representation [27, 48] with a unified latent feature capable of distilling diverse 2D foundation features, achieving both flexibility and efficiency. However, directly lifting 2D feature maps to dense, dynamic 3D Gaussians incurs significant computational costs and fails to scale with the spatial and temporal complexity of the scene. Inspired by recent advancements in Gaussian-based 4D reconstruction [34, 79], we propose to leverage the smooth, compact nature of underlying scene semantics and represent the dense 4D fea-

ture field using a sparse set of base or “Scaffold” features, enabling efficient representation and scalable adaptation to various downstream tasks.

Our pipeline is fully end-to-end differentiable, supervised by ground truth color and feature maps exported from 2D vision foundation models. Its general and model-agnostic design supports a wide range of vision tasks, spanning 2D (semantic and promptable segmentation), 3D (scene editing), and 4D (spatial and temporal VQA). Furthermore, our compact and versatile 4D feature field serves as a bridge, seamlessly integrating these tasks with LLMs in a continuous feedback loop. This integration enables intuitive and efficient execution through free-form, high-level natural language inputs or direct user interactions. By unifying diverse vision tasks, our framework—seamlessly integrated with LLMs—paves the way for the development of advanced, contextually and spatiotemporally aware 4D agentic AI systems.

We summarize our **contributions** as follows:

- We propose a general 4D feature field distillation technique using Gaussian Splatting to build an interactive 4D scene and lift the functionalities of any 2D vision (image/video) foundation models into the 4D realm, relying solely on monocular video input.
- We introduce the first compact and versatile 4D Gaussian feature field representation, which leverages the smooth, low-rank nature of scene semantics by modeling the dense 4D feature field as interpolations of a sparse set of base features.
- We develop an LLM-powered agentic AI capable of interpreting natural language prompts, dynamically adjusting configuration parameters, and iteratively refining results through trial and feedback, enabling intelligent 4D scene interaction and manipulation.

## 2. Related Works

### 2.1. 4D Representations for Reconstruction and Generation

Various 4D representations are designed to support a wide variety of 4D tasks. Among them, 4D reconstruction is the most widely studied topic that drives the field. Early methods [7, 11, 15, 17, 18, 32, 39, 54, 56–58, 62, 73, 82] implement variants of Neural Radiance Fields [51] to represent a dynamic 4D scene through implicit or hybrid representations. More recently, 3D Gaussian [27] has prevailed as a popular explicit 3D scene representation [9, 92, 99] and various variants [12, 22, 23, 34, 42, 48, 49, 72, 79, 83, 89, 90] are similarly proposed to represent dynamic 4D scenes. Additionally, the field of 4D Generation [44, 67, 68, 86, 91] has witnessed great advancements in representations that are more suitable for generative settings. In summary, most 4D representations can be categorized into deformation-based

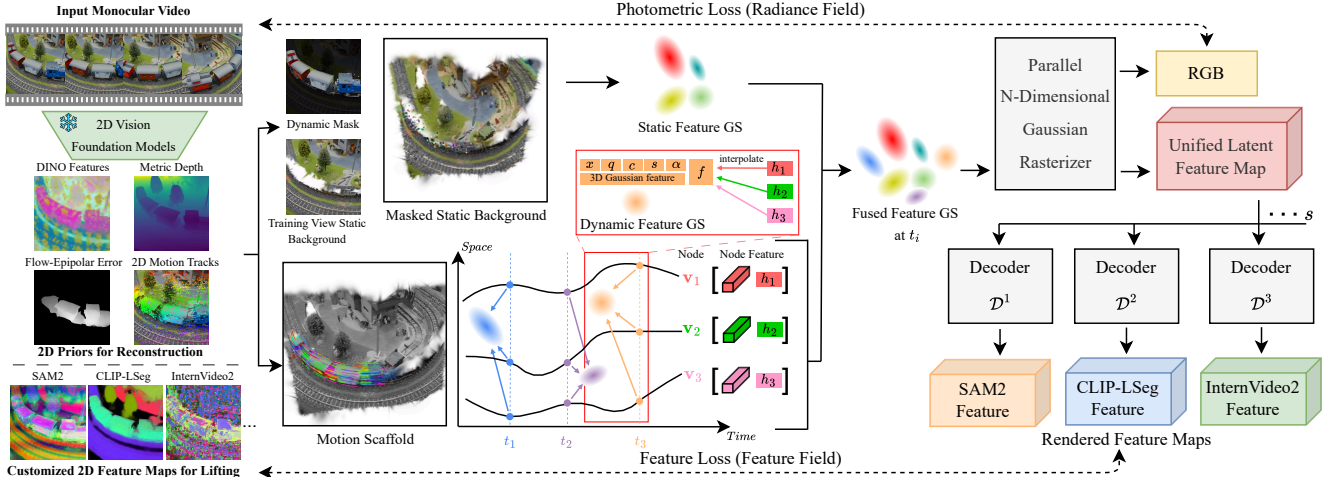


Figure 2. **Method overview.** Given an input monocular video, we infer 2D priors to segment static background (represented by static 3D Gaussians augmented with latent features) and dynamic foreground (represented by dynamic 3D Gaussians guided by Motion Scaffolds [34], a set of nodes  $\{v_i\}$  encoding 3D motion trajectories and latent features  $h_i$ ). Dynamic Gaussian features and motions are computed via interpolation from their  $K$ -nearest scaffold nodes. At each timestep, dynamic Gaussians are warped and fused with static Gaussians. A parallel rasterization [98] generates RGB images and a unified latent feature map, decoded into task-specific features—illustrated here by SAM2 [66], CLIP-LSeg [35], and InternVideo2 [81] for representative 2D (novel view segmentation), 3D (scene editing), and 4D (spatiotemporal VQA) tasks. Our framework generalizes to any 2D vision foundation model and is trained end-to-end using input RGB frames and customized features from pretrained 2D models. At inference, rendered feature maps from arbitrary views and timesteps are directly fed into task-specific decoders, seamlessly supporting user prompts and LLM interactions to form a unified 4D agentic AI system.

approaches [41, 48, 55, 56, 61, 74, 77, 84], temporal extended approaches [7, 12, 16, 89, 95], or an ensemble of multiple time-varying representations [19, 38, 40, 68, 85]. In our work, we extend the direction of dynamic 3D Gaussian Splatting to a versatile Gaussian feature field that simultaneously embraces multiple features for various vision tasks beyond 4D reconstruction and generation. To the best of our knowledge, this direction has not been widely explored.

## 2.2. 4D Reconstruction from Monocular Video

In the context of 4D reconstruction tasks, most methods [43, 53] focus on reconstructing dynamic 4D scenes from multiple well-calibrated videos. In contrast, another popular direction is to synthesize generic 4D scenes from monocular videos [19, 40, 55, 56, 61, 77, 84, 85]. For unposed casual videos, non-rigid structure-from-motion workflows [3–5, 10, 20, 24, 25, 31, 36, 52, 65, 76, 87, 88, 96] are adopted for estimating camera poses for each frame for further processing. Shape-of-motion [79] and MoSca [34] develop end-to-end workflows that reconstruct 4D scenes from casual captured videos with the help of 3D Gaussians [27, 47, 48]. Another series of work [38, 67, 68, 91] assumes the camera is fixed for the video and reconstructs 3D in the camera coordinate space. More recently, the success of DUST3R [80], a transformer-based novel paradigm for 3D reconstruction of arbitrary image collections, has been adapted for dynamic scenes in MonST3R [94]. In

comparison, we aim to perform monocular 4D reconstruction and feature lifting simultaneously, which cannot be achieved in previous approaches.

## 2.3. Feature Field Distillation

Research in novel view synthesis and feature field representation has been extensively developed within the NeRF framework [50] and 3D Gaussians [27]. Seminal works such as Semantic NeRF [97] and Panoptic Lifting [71] have integrated semantic information from segmentation networks into 3D spaces, revealing that combining noisy or inconsistent 2D labels within a 3D context can produce clear and accurate 3D segmentations. Building on this concept, methods such as those in [69] have shown that minimal input, like basic foreground-background masks, can be effective for object segmentation in 3D. Moving beyond label estimation to optimize NeRF, approaches like Distilled Feature Fields [30], NeRF-SOS [13], LERF [28], and Neural Feature Fusion Fields [78] have embedded pixel-aligned feature vectors from tools such as LSeg or DINO [8] into NeRF structures. More recently, numerous works [33, 45, 63, 70, 93, 98, 100] adopt similar strategies to distill information from well-trained 2D models to 3D Gaussians. However, feature lifting into 4D fields is not properly solved yet. In this work, we make the first attempt to deliver a versatile 4D Gaussian feature field framework that embraces multiple task features simultaneously.

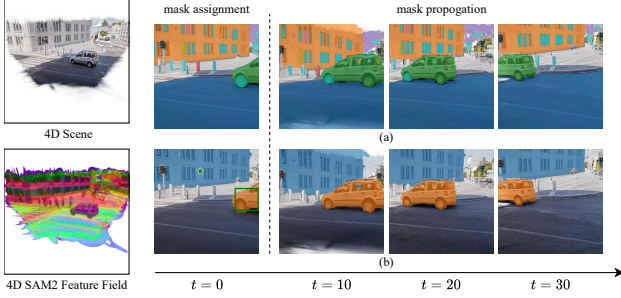


Figure 3. **Segment Anything in Dynamic 4D Scenes with SAM2 Feature Field.** For any rendered novel view video, we support: (a) Promptless segmentation (segment everything): when no user prompt is provided, segmentation masks are automatically assigned at the first frame ( $t = 0$ ) and then propagated across all frames. (b) Promptable segmentation (segment anything): the user can segment any object—static or dynamic—at any timestep using a point or box prompt, and the corresponding mask is robustly tracked and propagated through subsequent frames.

### 3. Method

Given a monocular RGB video  $\mathcal{I} = I_1, \dots, I_t$ , we reconstruct the underlying 4D scene as a set of dynamic 3D Gaussians (Sec. 3.1), each augmented with a unified latent feature embedding that distills informative features from various 2D foundation models for downstream tasks (Sec. 3.2). To address the challenge of handling high-dimensional features across a large number of Gaussians and timesteps, we adopt a compact feature representation (Sec. 3.3).

#### 3.1. Preliminaries: Dynamic 3D Gaussian Splatting

Following state-of-the-art monocular video-based dynamic 4D scene reconstruction approaches [34, 72, 79], we represent the scene with dynamic 3D Gaussians [48], namely a set of persistent 3D Gaussians [27] that deform over time  $t$ . Specifically, we base our approach on MoSca [34], which can reconstruct the dynamic 4D scene from monocular casual video. Like other methods, MoSca addresses the challenge from single-view partial observation by leveraging priors from 2D foundation models [2, 21, 26, 60, 75] and by regularizing Gaussian motion trajectories. At the core of MoSca’s method is a structured graph named 4D Motion Scaffold ( $\mathcal{V}, \mathcal{E}$ ) that drives the deformation of individual 3D Gaussians  $\mathcal{G} = \{G_j\}_{j=1}^n$  (see Fig. 2). Each node  $\mathbf{v}^{(i)} \in \mathcal{V}$  describes a 3D motion trajectory  $[\mathbf{Q}_1^{(i)}, \dots, \mathbf{Q}_t^{(i)}]$ ,  $\mathbf{Q} = [\mathbf{R}, \mathbf{t}] \in SE(3)$  with a control radius  $r^{(i)}$ . Edges  $\mathcal{E}$  describe a  $k$ -nearest neighbor graph over the motion trajectories. Given any 3D Gaussian located at  $\mathbf{x}$  at time  $\tau$ , to compute its deformation to time  $\tau'$ , we first find the trajectory node  $\mathbf{v}^{(i^*)}$  with the closest position  $\mathbf{p}_\tau^{(i)}$  at time  $\tau$ , i.e.,  $i^* = \arg \min_i \|\mathbf{p}_\tau^{(i)} - \mathbf{x}\|$ , then interpolate the deformation from  $K$  nearest trajectory nodes  $\{\mathbf{v}^{(i)}\}_{i \in \mathcal{E}(i^*)}$ , with weights  $\{w_i\}$  computed from the Gaussian-to-trajectory distance

and the node control radius. In other words, the dense, per-Gaussian motion trajectories are interpolations of a much smaller ( $100\times$ ) set of trajectory nodes. This design leverages the low-rank nature of the underlying scene motion and effectively regularizes Gaussian motions under limited supervision. We follow [34]’s protocols to initialize and train 4D Motion Scaffolds and dynamic Gaussians that represent the dynamic foreground elements, as well as another set of static Gaussians that model the static background. Please refer to our supplementary material for full details.

#### 3.2. Unified Latent Feature Distillation

We aim to go beyond appearance reconstruction by building a unified, versatile, and dynamic 4D feature field representation, capable of supporting diverse downstream visual tasks  $\mathcal{T} = \{T^1, \dots, T^S\}$  spanning 2D, 3D, and 4D, such as novel-view segmentation (2D), scene editing (3D), and scene-level spatiotemporal VQA (4D). A straightforward way to construct this 4D feature field is to extend existing 3D feature field frameworks [98] by replacing their 3D reconstruction modules with dynamic 4D reconstruction. However, this approach addresses each downstream task  $T^s$  independently: input frames  $\mathcal{I} = \{I_1, \dots, I_t\}$  are encoded into task-specific 2D feature maps  $\mathcal{F}^s = \{F_1^s, \dots, F_t^s\}$  using dedicated vision foundation model encoders  $E^s$ , leading to separate feature fields per task. Consequently, performing three tasks would require three separate reconstructions, making this strategy inefficient and redundant.

To overcome this limitation, we propose distilling a unified 4D feature field to coherently fuse diverse 2D vision foundation model features across views and timesteps, providing consistent feature access in 4D (3D space + time) for various downstream tasks. Specifically, as shown in Fig. 2, we utilize the parallel N-dimensional Gaussian rasterizer [98] to simultaneously render RGB images along with a unified latent feature  $\mathcal{F}$ , which is jointly learned with a set of lightweight decoders  $\{\mathcal{D}^1, \dots, \mathcal{D}^S\}$ . Each decoder maps the shared latent feature  $\mathcal{F} \in \mathbb{R}^D$  into task-specific target features  $\mathcal{F}^s \in \mathbb{R}^{D_s}$ , where  $D$  and  $D_s$  denote the dimensions of the unified latent and task-specific features, respectively, with  $D \ll D_s$ . Compared to separately optimizing high-dimensional features for each task, our unified latent representation significantly reduces computational overhead during rasterization. During the optimization process, we attach the feature vector  $f_j \in \mathbb{R}^D$  to each 3D Gaussian  $G_j \in \mathcal{G}$ , warp  $G_j$  to the target timestep  $\tau$  following the process introduced in [34], and rasterize  $f_j$  the same way we rasterize Gaussian color  $c_j$  as [98]. Conceptually, the RGB and feature reconstruction from viewpoint  $v$  at timestep  $\tau$  are computed as:

$$\begin{aligned} \hat{I}_\tau^v &= \text{Rasterize}(v, \{\text{warp}(G_j, \tau), c_j\}_{G_j \in \mathcal{G}}) \\ \hat{F}_\tau^v &= \text{Rasterize}(v, \{\text{warp}(G_j, \tau), f_j\}_{G_j \in \mathcal{G}}) \end{aligned} \quad (1)$$

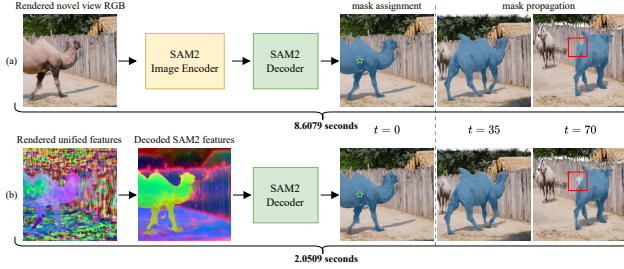


Figure 4. **Baseline Comparison on SAM2 Inference.** We compare segmentation quality and inference speed between (a) the naive RGB-based approach and (b) our feature-based method. Ours achieves comparable segmentation, accurately tracking the object over time, and avoids RGB artifacts (red box region at  $t = 70$ ), while reducing inference time to about  $4\times$  speed-up.

The resulting feature map  $\hat{F}_\tau$  (with viewpoint  $\mathbf{v}$  omitted for simplicity) is decoded by  $\mathcal{D}^s$  into a task-specific feature map  $\hat{F}_\tau^s$ , supervised by the ground-truth feature map obtained from a customized vision model encoder  $E^s$  (e.g., SAM2). Concretely, we optimize the following feature loss  $L_{\text{feat}}$  to jointly learn the feature field, alongside the photometric loss from [34] used for learning the radiance field.

$$L_{\text{feat}} = \sum_{s=1}^S \text{MSE}(\hat{F}_\tau^s, F_\tau^s), \quad (2)$$

$$\hat{F}_\tau^s = \mathcal{D}^s(\hat{F}_\tau), \quad F_\tau^s = E^s(I_\tau). \quad (3)$$

Note that the feature loss and photometric loss are independent, and introducing the feature field does not degrade the quality of the radiance field reconstruction (see Tab. 1).

### 3.3. Scaffold-Based Compact Feature

While our unified feature field effectively reduces feature dimensionality, optimizing a feature vector for every Gaussian remains computationally expensive. However, semantic features tend to be smooth in 3D and exhibit strong local correlations, similar to Gaussian motion trajectories. To leverage this correlation, we propose representing per-Gaussian features  $\{f_j\}$  as linear combinations of a smaller ( $100\times$  fewer) and more compact set of base features  $\{h_i\}$ , conveniently attached to nodes  $\{\mathbf{v}^{(i)}\}$  of the 4D Motion Scaffold (see Fig. 2). Recall that we compute per-Gaussian deformations by interpolating from their  $K$ -nearest trajectory nodes  $\{\mathbf{v}^{(i)}\}_{i \in \mathcal{E}(i^*)}$  with interpolation weights  $\{w_i\}$ . We reuse these same weights to obtain the per-Gaussian unified features as:

$$f_j = \sum_{i \in \mathcal{E}(i^*)} w_i h_i. \quad (4)$$

This compact feature representation significantly reduces the number of parameters required for optimization and

provides structural regularization, encouraging smoother learned features.

### 3.4. Interaction with AI Agent via Feature Fields

Agentic AI typically requires cross-modal interactions, notably between language and vision. In this work, we aim to build an AI agent that supports direct language interactions with our dynamic 4D scene representation. Such interaction requires a shared feature space between text and Gaussian features. However, text features are usually high-dimensional. For instance, language-guided scene editing requires CLIP features of 512 dimensions. Directly assigning these high-dimensional features to each Gaussian is computationally expensive. Feature 3DGS [98] addresses this by assigning each Gaussian a full 512-dimensional CLIP feature via a parallel  $N$ -dimensional Gaussian rasterizer, enabling direct language interaction but resulting in slow rendering speeds and excessive memory usage. Although Feature 3DGS proposes a CNN-based acceleration, it operates only on 2D rendered maps, failing to directly resolve the feature mismatch in 3D. In contrast, we optimize a compact, lower-dimensional Gaussian feature ( $D = 32$ ), train an MLP-based decoder on rendered 2D features, and apply it directly to 3D Gaussian features during inference. Given the intrinsic flexibility of MLPs, our approach efficiently bridges language features and our compact feature fields, enabling direct interaction between LLMs and our 4D scene through language. Furthermore, by leveraging InternVideo2 features, which can be rendered and decoded from any viewpoint in 3D space over time, we naturally lift the video chatbot LLM from 2D to 4D, enabling free-form language interaction with the AI agent within 4D space.

Beyond direct language interaction, LLM agents can operate in the loop for manipulation tasks by interpreting user prompts, automatically optimizing hyperparameters, and iteratively refining results for downstream applications. For example, given a scene-editing prompt like “Delete the dog,” the agent parses the operation command (“delete”) and the target object (“dog”), then generates an editing configuration with relevant hyperparameters, such as the softmax threshold for matching Gaussians to the target object. It tests various thresholds, evaluates the quality of rendered image samples, and selects the optimal configuration. This is then consistently applied across the entire 4D scene, enabling efficient, intelligent editing with minimal user input.

This perception-reasoning-action loop empowers the LLM-driven 4D agent to interpret, execute, and refine complex scene manipulations, making it a powerful tool for adaptive 4D scene editing and interaction. The system is particularly valuable for applications like interactive VR content creation and editing, where dynamic scene understanding and precise, context-aware modifications are essential.




Figure 5. **Semantic 4D Scene Understanding with CLIP Feature Field.** By lifting CLIP-LSeg [35] features into a 4D feature field, we enable pixel-level semantic segmentation from any view at any timestep. This allows robust 4D scene understanding, even as object appearances change over time—for example, accurately identifying a blooming flower from bud to full bloom across views.

Method	PSNR $\uparrow$	mIoU $\uparrow$	accuracy $\uparrow$	Size (MB) $\downarrow$
MoSca	25.166	-	-	<b>67.726</b>
MoSca + Feature 3DGS	25.191	0.506	<b>0.881</b>	593.907
Ours (single CLIP head)	25.186	<b>0.510</b>	0.880	95.294
Ours (full model)	<b>25.197</b>	0.503	0.876	95.457

Table 1. **Semantic segmentation on the Nvidia dataset [46].** Our method achieves comparable radiance reconstruction (PSNR) and segmentation performance, while significantly reducing memory usage compared to the baselines.

## 4. Experiments

In this section, we present experimental results and analyses of the capabilities of our proposed Feature4X framework. Based on the type of desired output, we categorize tasks into 2D (segmentation; see Secs. 4.1 and 4.2), 3D (scene editing; see Sec. 4.3), and 4D (spatiotemporal VQA; see Sec. 4.4).

### 4.1. Scene Interaction with Segment Anything

Segment Anything Model 2 (SAM2) [66] is an advanced segmentation model for promptable visual segmentation across images and videos, supporting various input types such as points and bounding boxes for interactive and precise results. In our approach, we extract per-frame features using SAM2’s image encoder and lift them into a 4D feature field, enabling direct decoding from novel-view rendered feature maps. This allows segmentation that is inherently aware of both viewpoint changes and temporal dynamics, as shown in Fig. 3. Masks are seamlessly propagated through space and time, regardless of object motion or camera trajectory, supporting both automatic and user-guided segmentation with temporal coherence.

Notably, our approach bypasses the need to first render RGB videos and then apply the full SAM2 pipeline (see (a) in Fig. 4). We observe that artifacts in novel view RGB rendering can mislead SAM2 during encoding, resulting in

ambiguous and inaccurate segmentation despite smoother masks. In contrast, our method achieves faster and more robust segmentation in 4D scenes by operating directly in feature space.

### 4.2. Scene Understanding with Semantics

To achieve pixel-level semantic segmentation across any novel view and timestep, we extend the capabilities of CLIP-LSeg [35] by lifting its features into a 4D CLIP feature field. CLIP-LSeg is a language-driven semantic segmentation model that leverages the Contrastive Language-Image Pre-training (CLIP) framework to align textual descriptions with visual content, enabling zero-shot segmentation by interpreting descriptive input labels. By lifting these 2D features into a 4D representation, our approach captures both spatial and temporal information, facilitating consistent semantic understanding in dynamic scenes, even when objects undergo significant appearance changes over time. As depicted in Fig. 5, our model accurately identifies a flower throughout its blooming process—from bud to full bloom—across various viewpoints and timesteps. This demonstrates the model’s proficiency in maintaining consistent semantic understanding despite substantial visual transformations in dynamic scenes.

Tab. 1 presents a quantitative comparison between our method and baselines on the Nvidia dataset [46]. “MoSca [34]” is the baseline 4D reconstruction method, while “MoSca + Feature 3DGS [98]” is a naive, non-compact baseline using RGB + feature rasterization. Our models—using a 32-dimension latent feature and 512-dimension CLIP-LSeg [35] as the target—maintain comparable PSNR to MoSca, showing no loss in reconstruction quality. The single CLIP head variant achieves the highest mIoU, demonstrating effective feature distillation. Although the naive baseline yields slightly higher accuracy, our models are about  $6.2\times$  more space-efficient, with the full model delivering competitive performance while remaining compact and general-purpose.

### 4.3. Scene Editing with Language and AI Agent

We employ an LLM agent (GPT-4o [1]) to interpret natural language prompts, optimize editing parameters, execute precise queries, and iteratively refine results for intelligent 3D scene editing. The agent begins by parsing the user’s high-level prompt to identify the editing action—such as “extract,” “delete,” or “change color”—and the target object (e.g., “cow”), as shown in Fig. 6. Based on the task, it generates candidate configurations with varying parameters relevant to the prompt. Specifically, we first compute the probability of each 3D Gaussian corresponding to a language description  $l$  of the target object as:

$$\mathbf{p}(l | j) = \frac{e^s}{\sum_{s_i \in \mathcal{L}} e^{s_i}}, \quad s = \frac{q(l) \cdot f_j}{\|q(l)\| \|f_j\|}, \quad (5)$$

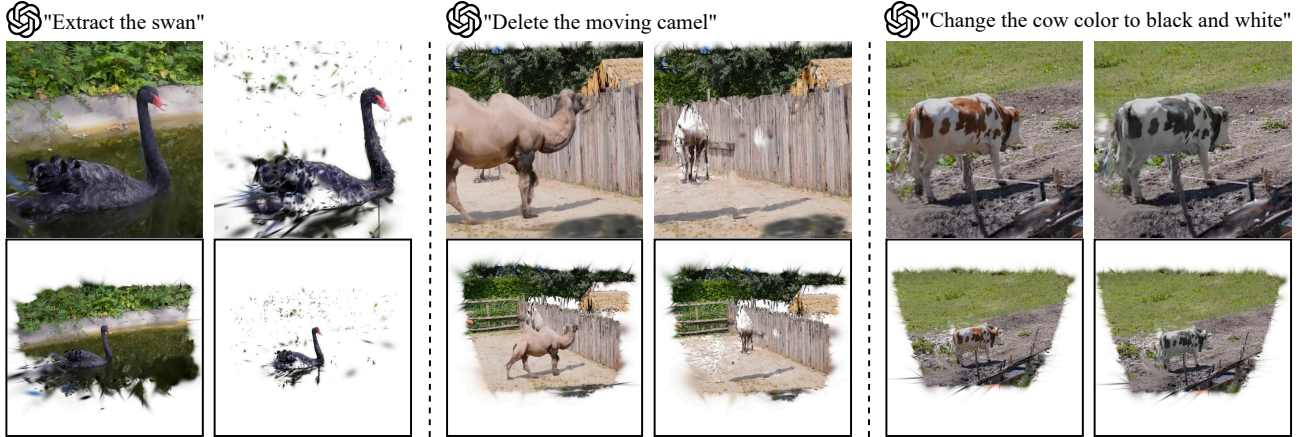


Figure 6. **Scene Editing with AI Agent.** Given user prompts, our GPT-powered agent interprets editing intent and autonomously performs scene edits via our 4D CLIP feature field. Examples include both geometric (e.g., “extract” and “delete”) and appearance (e.g., “change color”) editing in 3D space. While results may not be perfect due to imperfect fine-grained feature alignment and non-optimal editing parameter tuning, the agent adaptively refines parameters and applies edits consistently across views and time—greatly reducing the need for manual tuning—and demonstrates robust, interactive 4D scene manipulation.

where  $q(l)$  is the text feature,  $f_j$  is the CLIP feature of Gaussian  $j$ , and  $\mathcal{L}$  is the label set of all possible scene objects. The agent filters Gaussians using various thresholds based on the softmax probability  $\mathbf{p}(l | j)$ , applies the edit configuration, and evaluates the results via rendered image samples. It selects the parameter setting that best aligns with the intended edit and applies it consistently across all video frames, ensuring coherent editing throughout the 4D scene.

This GPT-powered 4D Agentic AI system interprets, executes, and refines complex scene manipulations, serving as an intelligent assistant for adaptive editing and interaction. By autonomously exploring, selecting, and tuning parameters—tasks typically handled by human editors—the agent enhances responsiveness and consistency. GPT acts as a self-refining decision-maker, illustrating the potential of agentic AI in creative and technical workflows where iterative optimization and goal-directed behavior significantly streamline and elevate editing tasks.

#### 4.4. Scene Reasoning with 4D Chatbot Agent

While large vision-language models (VLMs) have demonstrated impressive performance in visual question answering (VQA) tasks, their application has predominantly been confined to 2D visual modalities, such as images and videos. To our knowledge, we are the first to extend the capabilities of a VLM (Video-LLM) into 4D—encompassing three spatial dimensions plus time—by integrating VLM knowledge into our 4D feature field. This advancement enables a more comprehensive understanding and interaction with dynamic 4D scenes, capturing both spatial and temporal nuances beyond the static 2D perspective.

We utilize InternVideo2 [37, 81]—a state-of-the-art video

foundation model—to extract comprehensive video features. Specifically, we sample the input video into several clips, each comprising 8 frames sampled at equal intervals to ensure enough GPU memory available during the whole process. These clips are processed through InternVideo2’s vision transformer encoder to obtain segment-level features and then aggregated together. Additionally, for each segment, InternVideo2 generates a class token representing high-level semantic information. We compute the average of these class tokens to form a global video representation. This class feature is then concatenated with the aggregated segment features, effectively integrating temporal information which is necessary for dynamic scene reasoning.

In our system, free-form visual question answering (VQA) is supported through a language model [37], enabling users to interact with a chatbot agent using natural language while exploring 4D dynamic scenes, as illustrated in Fig. 7. Unlike prior approaches that rely solely on 2D monocular video—limited to a single view at a single time step—our method leverages the reconstructed 4D radiance and feature fields to support richer inference sources, including both local novel views (via free camera movement) and global novel views (via zoom-out for broader context). In the example shown in Fig. 7, we observe that InternVideo2, when operating on the input monocular video, fails to correctly answer spatial and temporal questions—even when the answers are obvious to humans—while the same model, using our global novel view feature, succeeds (see Tab. 2). This is because monocular video lacks the spatial coverage and contextual cues necessary for spatiotemporal reasoning. In contrast, our framework allows rendering of novel views and their cor-

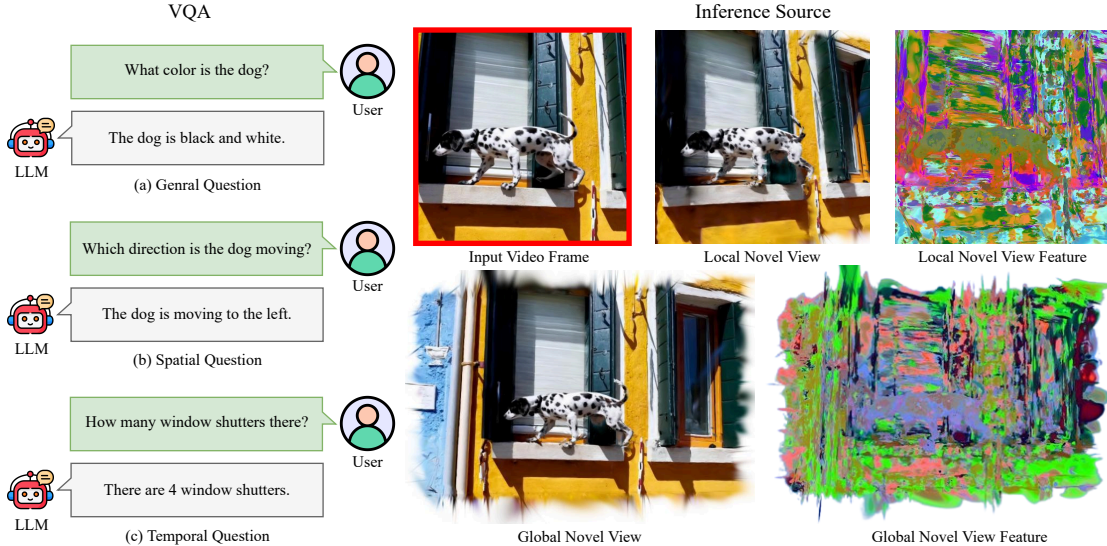


Figure 7. **VQA with Chatbot Agent.** (Left) Our model supports free-form VQA across diverse question types—general, spatial, and temporal—by distilling InternVideo2 [81] features. (Right) At each timestep, we reconstruct both a 4D radiance field and a 4D feature field, providing more inference sources beyond the input video frame—including local (moving camera) and global (zoomed-out) novel views and their corresponding feature maps—thereby supporting VQA in 4D and enhancing the model’s spatiotemporal reasoning capabilities.

Inference Source	General	Spatial	Temporal	Time (s)
Input Video View	✓	×	×	33.65
Local Novel View	✓	✓	×	33.86
Local Novel Feature	✓	×	✓	11.99
Global Novel View	✓	×	✓	33.37
Global Novel Feature	✓	✓	✓	12.24

Table 2. **VQA performance across different inference sources with the scene shown in Fig. 7.** Feature-based inference supports all question types with lower latency, while view-based methods are limited in spatial and temporal reasoning.

Inference Source	Spatial Acc.	Temporal Acc.	Overall Acc.	Time (s)
Input Video View	48.50	49.84	49.06	10.02
Local Novel View	49.75	51.13	50.31	9.78
Local Novel Feature	54.50	54.37	56.29	<b>2.81</b>
Global Novel View	47.00	49.19	47.48	10.06
Global Novel Feature	<b>58.75</b>	<b>58.25</b>	<b>61.32</b>	3.42

Table 3. **Spatiotemporal VQA on DAVIS dataset [59].** Compared to 2D video inference, our 4D feature space inference (Global Novel Feature) enhances the Video-LLM’s spatiotemporal reasoning while achieving approximately  $3\times$  faster inference.

responding feature maps at arbitrary time steps, injecting explicit signals that enhance reasoning. To evaluate this more systematically, we construct 400 objective questions across 50 scenes in the DAVIS dataset [59], covering spatial only (e.g., “Is the person facing left?”), temporal only (e.g., “How many times does the dancer spin?”), and spatiotemporal (e.g., “Which direction is the duck moving?”) categories. As shown in Tab. 3, our method improves VQA accuracy by enriching scene understanding through 4D feature field reconstruction, while also accelerating inference by eliminating the need for video encoding.

## 5. Conclusion

We have introduced Feature4X, a general and scalable framework that bridges casually captured monocular videos to interactive 4D agentic AI systems. By distilling diverse 2D vision foundation model features into a unified, dynamic 4D Gaussian feature field, our method supports a wide range of tasks across 2D (segmentation), 3D (scene editing), and 4D (spatiotemporal VQA). Through a compact and versatile representation, Feature4X achieves efficient training and inference while maintaining high-quality appearance reconstruction and robust semantic understanding. Furthermore, by integrating LLMs for language-based interaction and autonomous task refinement, our system elevates traditional perception tasks into a perception-reasoning-action loop, enabling intelligent and adaptive manipulation of dynamic 4D scenes. We hope Feature4X opens new avenues in 4D vision and agentic AI research by facilitating immersive, multimodal interaction with dynamic visual content.

**Acknowledgement** We especially thank Jiahui Lei for valuable guidance on MoSca implementation and insightful suggestions on experiments. This project was supported by LUCI program under the Basic Research Office and partially supported by ARL grants W911NF20-2-0158 and W911NF-21-2-0104 under the cooperative A2I2 program. Z.W. is supported by ARL and LUCI program, and an Army Young Investigator Award. L.G. is supported by a Vannevar Bush Faculty Fellowship. A.K. is supported by a DARPA Young Faculty Award, NSF CAREER Award IIS-2046737, and Army Young Investigator Award.



## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 6
- [2] Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023. 4
- [3] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *CVPR*, 1999. 3
- [4] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *ECCV*, 2016.
- [5] Aljaz Bozic, Pablo Palafox, Michael Zollöfer, Angela Dai, Justus Thies, and Matthias Nießner. Neural non-rigid tracking. In *NeurIPS*, 2020. 3
- [6] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. *arXiv preprint arXiv:2211.09800*, 2022. 2
- [7] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 130–141, 2023. 2, 3
- [8] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 3
- [9] Jaeyoung Chung, Suyoung Lee, Hyeongjin Nam, Jaerin Lee, and Kyoung Mu Lee. Luciddreamer: Domain-free generation of 3d gaussian splatting scenes. *arXiv preprint arXiv:2311.13384*, 2023. 2
- [10] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996. 3
- [11] Yilun Du, Yinan Zhang, Hong-Xing Yu, Joshua B Tenenbaum, and Jiajun Wu. Neural radiance flow for 4d view synthesis and video processing. In *ICCV*, 2021. 2
- [12] Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 4d gaussian splatting: Towards efficient novel view synthesis for dynamic scenes. *arXiv preprint arXiv:2402.03307*, 2024. 2, 3
- [13] Zhiwen Fan, Peihao Wang, Yifan Jiang, Xinyu Gong, De-jia Xu, and Zhangyang Wang. Nerf-sos: Any-view self-supervised object segmentation on complex scenes. *arXiv preprint arXiv:2209.08776*, 2022. 3
- [14] Zhiwen Fan, Jian Zhang, Wenyan Cong, Peihao Wang, Renjie Li, Kairun Wen, Shijie Zhou, Achuta Kadambi, Zhangyang Wang, Danfei Xu, et al. Large spatial model: End-to-end unposed images to semantic 3d. *Advances in Neural Information Processing Systems*, 37:40212–40229, 2024. 2
- [15] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12479–12488, 2023. 2
- [16] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12479–12488, 2023. 3
- [17] Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8649–8658, 2021. 2
- [18] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *CVPR*, 2021. 2
- [19] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, 2021. 3
- [20] Wei Gao and Russ Tedrake. Surfelpwarp: Efficient non-volumetric single view dynamic reconstruction. *arXiv preprint arXiv:1904.13073*, 2019. 3
- [21] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *European Conference on Computer Vision*, pages 59–75. Springer, 2022. 4
- [22] Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4220–4230, 2024. 2
- [23] Shengxiang Ji, Guanjun Wu, Jiemin Fang, Jiazhong Cen, Taoran Yi, Wenyu Liu, Qi Tian, and Xinggang Wang. Segment any 4d gaussians. *arXiv preprint arXiv:2407.04504*, 2024. 2
- [24] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 3
- [25] Abhishek Kar, Shubham Tulsiani, Joao Carreira, and Jitendra Malik. Category-specific object reconstruction from a single image. In *CVPR*, 2015. 3
- [26] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Co-tracker: It is better to track together. *arXiv preprint arXiv:2307.07635*, 2023. 4
- [27] 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, 3, 4
- [28] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lerp: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. 2, 3

- [29] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2
- [30] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. In *NeurIPS*, 2022. 2, 3
- [31] Nilesh Kulkarni, Abhinav Gupta, David Fouhey, and Shubham Tulsiani. Articulation-aware canonical surface mapping. In *CVPR*, 2020. 3
- [32] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *NeurIPS*, 2021. 2
- [33] Hyunjee Lee, Youngsik Yun, Jeongmin Bae, Seoha Kim, and Youngjung Uh. Rethinking open-vocabulary segmentation of radiance fields in 3d space. *arXiv preprint arXiv:2408.07416*, 2024. 2, 3
- [34] Jiahui Lei, Yijia Weng, Adam Harley, Leonidas Guibas, and Kostas Daniilidis. Mosca: Dynamic gaussian fusion from casual videos via 4d motion scaffolds. *arXiv preprint arXiv:2405.17421*, 2024. 2, 3, 4, 5, 6
- [35] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022. 3, 6
- [36] Hao Li, Robert W Sumner, and Mark Pauly. Global correspondence optimization for non-rigid registration of depth scans. In *Computer graphics forum*, pages 1421–1430. Wiley Online Library, 2008. 3
- [37] KunChang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. 2, 7
- [38] Renjie Li, Panwang Pan, Bangbang Yang, Dejie Xu, Shijie Zhou, Xuanyang Zhang, Zeming Li, Achuta Kadambi, Zhangyang Wang, and Zhiwen Fan. 4k4dgen: Panoramic 4d generation at 4k resolution. *arXiv preprint arXiv:2406.13527*, 2024. 3
- [39] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, and Zhaoyang Lv. Neural 3d video synthesis from multi-view video. In *CVPR*, 2022. 2
- [40] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 3
- [41] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8508–8520, 2024. 3
- [42] Yiqing Liang, Numair Khan, Zhengqin Li, Thu Nguyen-Phuoc, Douglas Lanman, James Tompkin, and Lei Xiao. Gaufré: Gaussian deformation fields for real-time dynamic novel view synthesis. *arXiv preprint arXiv:2312.11458*, 2023. 2
- [43] Youtian Lin, Zuozhuo Dai, Siyu Zhu, and Yao Yao. Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21136–21145, 2024. 3
- [44] Huan Ling, Seung Wook Kim, Antonio Torralba, Sanja Fidler, and Karsten Kreis. Align your gaussians: Text-to-4d with dynamic 3d gaussians and composed diffusion models. *arXiv preprint arXiv:2312.13763*, 2023. 2
- [45] Xinyi Liu, Tianyi Zhang, Matthew Johnson-Roberson, and Weiming Zhi. Splatraj: Camera trajectory generation with semantic gaussian splatting. *arXiv preprint arXiv:2410.06014*, 2024. 2, 3
- [46] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 6
- [47] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20654–20664, 2024. 3
- [48] 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, 3, 4
- [49] Marko Mihajlovic, Sergey Prokudin, Siyu Tang, Robert Maier, Federica Bogo, Tony Tung, and Edmond Boyer. Splatfields: Neural gaussian splats for sparse 3d and 4d reconstruction. In *European Conference on Computer Vision*, pages 313–332. Springer, 2025. 2
- [50] 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. *Communications of the ACM*, 65(1):99–106, 2021. 3
- [51] 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. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [52] Richard Newcombe, Dieter Fox, and Steve Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *CVPR*, 2015. 3
- [53] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5379–5389, 2019. 3
- [54] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *ICCV*, 2021. 2
- [55] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 3

- [56] 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)*, 40(6):1–12, 2021. 2, 3
- [57] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14314–14323, 2021.
- [58] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 2
- [59] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016. 8
- [60] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Matia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. *arXiv preprint arXiv:2403.18913*, 2024. 4
- [61] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *CVPR*, 2021. 3
- [62] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 2
- [63] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20051–20060, 2024. 2, 3
- [64] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [65] Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Reconstructing 3d human pose from 2d image landmarks. In *ECCV*, 2012. 3
- [66] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 2, 3, 6
- [67] Jiawei Ren, Liang Pan, Jiaxiang Tang, Chi Zhang, Ang Cao, Gang Zeng, and Ziwei Liu. Dreamgaussian4d: Generative 4d gaussian splatting. *arXiv preprint arXiv:2312.17142*, 2023. 2, 3
- [68] Jiawei Ren, Kevin Xie, Ashkan Mirzaei, Hanxue Liang, Xiaohui Zeng, Karsten Kreis, Ziwei Liu, Antonio Torralba, Sanja Fidler, Seung Wook Kim, et al. L4gm: Large 4d gaussian reconstruction model. *arXiv preprint arXiv:2406.10324*, 2024. 2, 3
- [69] Zhongzheng Ren, Aseem Agarwala<sup>†</sup>, Bryan Russell<sup>†</sup>, Alexander G. Schwing<sup>†</sup>, and Oliver Wang<sup>†</sup>. Neural volumetric object selection. In *CVPR*, 2022. (<sup>†</sup> alphabetic ordering). 3
- [70] Jin-Chuan Shi, Miao Wang, Hao-Bin Duan, and Shao-Hua Guan. Language embedded 3d gaussians for open-vocabulary scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5333–5343, 2024. 2, 3
- [71] Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Buló, Norman Müller, Matthias Nießner, Angela Dai, and Peter Kotschieder. Panoptic lifting for 3d scene understanding with neural fields. *arXiv preprint arXiv:2212.09802*, 2022. 3
- [72] Colton Stearns, Adam Harley, Mikaela Uy, Florian Dubost, Federico Tombari, Gordon Wetzstein, and Leonidas Guibas. Dynamic gaussian marbles for novel view synthesis of casual monocular videos. *arXiv preprint arXiv:2406.18717*, 2024. 2, 4
- [73] Shih-Yang Su, Frank Yu, Michael Zollhoefer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. In *NeurIPS*, 2021. 2
- [74] Jiakai Sun, Han Jiao, Guangyuan Li, Zhanjie Zhang, Lei Zhao, and Wei Xing. 3dstream: On-the-fly training of 3d gaussians for efficient streaming of photo-realistic free-viewpoint videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20675–20685, 2024. 3
- [75] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 4
- [76] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34:16558–16569, 2021. 3
- [77] 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 *ICCV*, 2021. 3
- [78] Vadim Tschernezki, Iro Laina, Diane Larlus, and Andrea Vedaldi. Neural Feature Fusion Fields: 3D distillation of self-supervised 2D image representations. In *3DV*, 2022. 2, 3
- [79] Qianqian Wang, Vickie Ye, Hang Gao, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4d reconstruction from a single video. *arXiv preprint arXiv:2407.13764*, 2024. 2, 3, 4
- [80] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 3

- [81] Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yanan He, Guo Chen, Baoqi Pei, Rongkun Zheng, Zun Wang, Yansong Shi, et al. Internvideo2: Scaling foundation models for multimodal video understanding. In *European Conference on Computer Vision*, pages 396–416. Springer, 2024. [2](#), [3](#), [7](#), [8](#)
- [82] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *CVPR*, 2022. [2](#)
- [83] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. [2](#)
- [84] Tianhao Wu, Fangcheng Zhong, Andrea Tagliasacchi, Forrester Cole, and Cengiz Oztireli. D<sup>2</sup>nerf: Self-supervised decoupling of dynamic and static objects from a monocular video. In *NeurIPS*, 2022. [3](#)
- [85] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *CVPR*, 2021. [3](#)
- [86] Dejia Xu, Hanwen Liang, Neel P Bhatt, Hezhen Hu, Hanxue Liang, Konstantinos N Plataniotis, and Zhangyang Wang. Comp4d: Llm-guided compositional 4d scene generation. *arXiv preprint arXiv:2403.16993*, 2024. [2](#)
- [87] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Huiwen Chang, Deva Ramanan, William T Freeman, and Ce Liu. Lasr: Learning articulated shape reconstruction from a monocular video. In *CVPR*, 2021. [3](#)
- [88] Gengshan Yang, Minh Vo, Neverova Natalia, Deva Ramanan, Vedaldi Andrea, and Joo Hanbyul. Banmo: Building animatable 3d neural models from many casual videos. In *CVPR*, 2022. [3](#)
- [89] 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](#), [3](#)
- [90] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20331–20341, 2024. [2](#)
- [91] Yuyang Yin, Dejia Xu, Zhangyang Wang, Yao Zhao, and Yunchao Wei. 4dgen: Grounded 4d content generation with spatial-temporal consistency. *arXiv preprint arXiv:2312.17225*, 2023. [2](#), [3](#)
- [92] Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T Freeman, and Jiajun Wu. Wonderworld: Interactive 3d scene generation from a single image. *arXiv preprint arXiv:2406.09394*, 2024. [2](#)
- [93] Justin Yu, Kush Hari, Kishore Srinivas, Karim El-Refai, Adam Rashid, Chung Min Kim, Justin Kerr, Richard Cheng, Muhammad Zubair Irshad, Ashwin Balakrishna, et al. Language-embedded gaussian splats (legs): Incrementally building room-scale representations with a mobile robot. *arXiv preprint arXiv:2409.18108*, 2024. [2](#), [3](#)
- [94] Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the presence of motion. *arXiv preprint arXiv:2410.03825*, 2024. [3](#)
- [95] Shuai Zhang, Huangxuan Zhao, Zhenghong Zhou, Guan-jun Wu, Chuansheng Zheng, Xinggang Wang, and Wenyu Liu. Togs: Gaussian splatting with temporal opacity offset for real-time 4d dsa rendering. *arXiv preprint arXiv:2403.19586*, 2024. [3](#)
- [96] Wang Zhao, Shaohui Liu, Hengkai Guo, Wenping Wang, and Yong-Jin Liu. Particlesfm: Exploiting dense point trajectories for localizing moving cameras in the wild. In *European Conference on Computer Vision*, pages 523–542. Springer, 2022. [3](#)
- [97] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew Davison. In-place scene labelling and understanding with implicit scene representation. In *ICCV*, 2021. [3](#)
- [98] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suyu You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21676–21685, 2024. [2](#), [3](#), [4](#), [5](#), [6](#)
- [99] Shijie Zhou, Zhiwen Fan, Dejia Xu, Haoran Chang, Pradyumna Chari, Tejas Bharadwaj, Suyu You, Zhangyang Wang, and Achuta Kadambi. Dreamscene360: Unconstrained text-to-3d scene generation with panoramic gaussian splatting. In *European Conference on Computer Vision*, pages 324–342. Springer, 2024. [2](#)
- [100] Xingxing Zuo, Pouya Samangouei, Yunwen Zhou, Yan Di, and Mingyang Li. Fmgs: Foundation model embedded 3d gaussian splatting for holistic 3d scene understanding. *International Journal of Computer Vision*, pages 1–17, 2024. [2](#), [3](#)