

Geo4D: Leveraging Video Generators for Geometric 4D Scene Reconstruction

Zeren Jiang¹ Chuanxia Zheng¹ Iro Laina¹ Diane Larlus² Andrea Vedaldi¹
¹Visual Geometry Group, University of Oxford ²Naver Labs Europe
{zeren, cxzheng, iro, vedaldi}@robots.ox.ac.uk diane.larlus@naverlabs.com
[geo4d.github.io](https://github.com/Geo4D)




Figure 1. Geo4D repurposes a video diffusion model [102] for monocular 4D reconstruction. It uses only synthetic data for training, yet generalizes well to *out-of-domain* real videos. It predicts several geometric modalities, including point maps, disparity maps, and ray maps, fusing and aligning them to obtain state-of-the-art dynamic reconstruction even for scenes with extreme object and camera motion.

Abstract

We introduce *Geo4D*, a method to repurpose video diffusion models for monocular 3D reconstruction of dynamic scenes. By leveraging the strong dynamic priors captured by large-scale pre-trained video models, *Geo4D* can be trained using only synthetic data while generalizing well to real data in a zero-shot manner. *Geo4D* predicts several complementary geometric modalities, namely point, disparity, and ray maps. We propose a new multi-modal alignment algorithm to align and fuse these modalities, as well as a sliding window approach at inference time, thus enabling robust and accurate 4D reconstruction of long videos. Extensive experiments across multiple benchmarks show that *Geo4D* significantly surpasses state-of-the-art video depth estimation methods.

1. Introduction

We consider the problem of *feed-forward 4D reconstruction*, which involves learning a neural network to reconstruct the 3D geometry of a dynamic scene from a monoc-

ular video. This task is particularly challenging for videos captured in uncontrolled settings, such as those shot with handheld cameras or downloaded from the Internet. However, a robust solution to this problem would have a tremendous impact on a wide range of applications, from video understanding to computer graphics and robotics.

4D reconstruction from videos is related to multi-view static 3D reconstruction, which is typically addressed using methods from visual geometry like bundle adjustment. Recent neural networks [89, 92] have emerged as powerful tools that can replace, or at least complement, bundle adjustment. They excel especially in difficult reconstruction scenarios, involving, *e.g.*, textureless surfaces and occlusions, thanks to the priors they learn from data. Given the additional challenges involved in 4D reconstruction, we expect that such priors would benefit this task even more.

In fact, powerful networks like DUS_t3R [92], designed for static multi-view 3D reconstruction, have recently been extended to the dynamic case, for example by MonST3R [113]. However, these models are heavily engineered to solve specific 3D reconstruction problems. Most importantly, they require significant amounts of training data with 3D annotations for supervision. Such data

is difficult to collect for dynamic scenes, especially in real life. This suggests using 4D synthetic training data instead. However, this data is difficult to obtain at scale, and the gap with the real world can compromise generalization.

One way to mitigate this problem is to pre-train the model on tasks related to 3D reconstruction for which real data is easily available. For example, DUS_t3R [92] and derived methods [113] use image matching for pre-training [98]. Here, we suggest starting instead from an off-the-shelf *video generator*. Video generators are powerful models, often considered proxies of world simulators [37, 54, 59]. More importantly for us, the videos they generate demonstrate an understanding of effects like camera motion and perspective, as well as typical object motion in the context of a scene. However, they only generate pixels, leaving any 3D or 4D understanding *implicit* and thus not directly actionable.

In this work, we show that a *pre-trained off-the-shelf video generator* can be turned into an effective *monocular feed-forward 4D reconstructor*. To this end, we introduce **Geo4D**, a novel approach for adapting Video Generators for **Geometric 4D Reconstruction**. With Geo4D, we demonstrate that these generic video architectures can successfully solve complex 4D reconstruction tasks, which is a step towards future video foundation models that natively integrate 4D geometry. Prior work such as Marigold [28] and concurrent work DepthCrafter [22] have looked at adapting, respectively, image and video generators for depth estimation. Here, we go one step further and consider *the full recovery of 4D geometry, including camera motion and dynamic 3D structure*.

With Geo4D, our goal is to make 4D geometry explicit in the video generator. This in turn requires us to choose an *explicit representation* of 4D information. We follow DUS_t3R and adopt its viewpoint-invariant point maps. Namely, we associate each pixel in each frame with the coordinate of the corresponding 3D point, expressed relative to the first frame in the video, used as a reference. Hence, the static parts of the point clouds extracted from the different frames line up, and the dynamic parts form a 3D ‘trace’ of the motion of the dynamic objects, as shown in Fig. 1.

Viewpoint-invariant point maps are a powerful representation because they implicitly encode the camera motion and intrinsics and can be easily predicted by a neural network [92]. However, they are not necessarily the best representation for all parts of the scene, particularly for points far away from the observer or even at infinity, such as the sky. We thus consider two more *modalities* with better dynamic range, namely disparity maps and camera ray maps. Ray maps, in particular, are defined for all image pixels regardless of the scene geometry.

Our model thus predicts three modalities: point, disparity, and ray maps. These modalities are redundant in prin-

ciple, but complementary in practice. At test time, we reconcile them via a fast, global optimization step and show that this leads to significantly more robust 4D reconstructions. Due to depth and ray map prediction, we show very strong empirical results on video depth estimation and in the recovery of the camera orientation.

One of the challenges of monocular 4D reconstruction is that it is ambiguous, significantly more so than static 3D reconstruction. However, the stochastic nature of the video generator can help deal with this ambiguity. We also introduce uncertainty maps in the encoder-decoder architecture that processes the geometric maps, and integrate them into the multi-modal alignment process.

Overall, our contributions are as follows. (i) We introduce Geo4D, a 4D feed-forward network for dynamic scene reconstruction that builds on top of an off-the-shelf video generator. (ii) We suggest generating multiple partially redundant geometric modalities and fusing them at test time via lightweight optimization. (iii) We show the benefits of this multi-modal fusion in terms of improved 4D prediction accuracy. Experiments show that this model can reconstruct even highly dynamic scenes (such as the drifting scene in DAVIS [23] presented in Fig. 1) and outperforms current video depth and camera rotation estimation methods.

2. Related Work

2.1. Dynamic Scene Reconstruction

Static 3D reconstruction. Feed-forward 3D reconstruction has achieved remarkable success across various representations, including voxels [11, 74, 83], meshes [18, 72, 90], and point clouds [41, 110]. These advancements have been further driven by implicit neural representations [52, 56, 60, 75] and the emergence of 3D Gaussian Splatting (3D-GS) [7, 9, 29, 76, 79, 80]. Recently, DUS_t3R [92] introduced a point map representation for scene-level 3D reconstruction, followed by [35, 86, 89, 104]. However, these models predominantly focus on *static* 3D reconstruction. Our approach also uses point maps as a representation but extends them to handle *dynamic* scenes, which present additional challenges due to object motion over time.

Iterative 4D reconstruction. Iterative or optimization-based approaches reconstruct 4D models from monocular videos by iteratively fitting the observed data. Classical techniques often rely on RGB-D sensors [24, 53], but such steps are impractical for many real-world scenes. Recently, with advancements in neural representations [52, 56], NeRF-based approaches [27, 38, 39, 57, 58, 62] have shown impressive results. However, volume rendering in NeRF is computationally expensive. Convergence and rendering speed can be improved by using 3D-GS representations [12, 29, 34, 43, 91, 99, 107, 111], which reduce but do not eliminate the cost of iterative optimization. Very

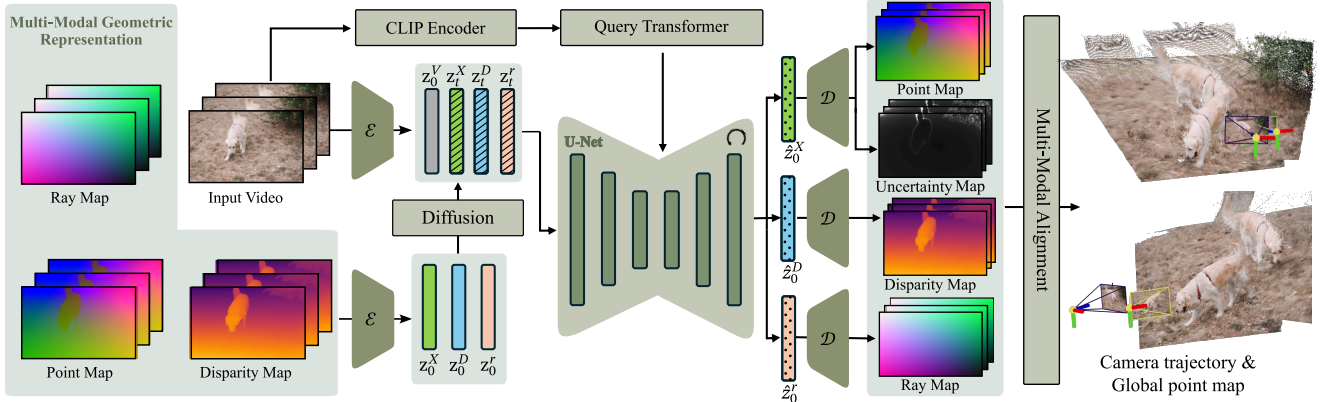


Figure 2. **Overview of Geo4D.** During training, video conditions are injected by locally concatenating the latent feature of the video with diffused geometric features z_t^X, z_t^D, z_t^R and are injected globally via cross-attention in the denoising U-Net, after CLIP encoding and a query transformer. The U-Net is fine-tuned via Eq. 2. During inference, iteratively denoised latent features $\hat{z}_0^X, \hat{z}_0^D, \hat{z}_0^R$ are decoded by the fine-tuned VAE decoder, followed by multi-modal alignment optimization for coherent 4D reconstruction.

recently, MegaSaM [40] achieved highly accurate and robust camera pose estimation and reconstruction for dynamic videos, but it requires accurate monocular depth priors. Similarly, Uni4D [108] produces accurate 4D reconstructions by leveraging various visual foundation models and performing multi-stage bundle adjustment. In contrast, our approach is a diffusion-driven feed-forward framework, which eliminates the need for per-video bundle adjustment and external depth estimation models.

Feed-forward 4D reconstruction. Similar to our approach, recent works have started to explore *feed-forward* 4D reconstruction for dynamic scenes: a monocular video with dynamic objects is processed by a neural network to recover a 4D representation. For objects, L4GM [66] and Animate3D [26] first generate multi-view videos from a monocular video input, and subsequently apply 3D-GS [29] to reconstruct a temporally consistent 4D model. For scenes, a notable example is MonST3R [113], which adapts the *static* scene reconstruction of DUS3R [92] to handle *dynamic* scenes. Very recently, Easi3R [8] applies attention adaptation during inference and performs 4D reconstruction based on DUS3R [92] in an efficient, training-free manner.

2.2. Geometric Diffusion Models

Our method builds upon advancements in video diffusion models [3, 4, 16, 19, 21, 31, 73, 88, 94, 102, 112], which generate *temporally consistent* videos from text or image prompts. Recent studies have explored the rich 3D priors embedded within large-scale pre-trained diffusion models, employing either knowledge distillation [25, 42, 51, 61, 87, 96] or fine-tuning [20, 36, 45–47, 71, 85, 118] for 3D reconstruction and generation. While these methods have significantly advanced single-object 3D reconstruction from sparse inputs, they remain largely constrained to *static, isolated* objects centered within an image. Beyond single-

object reconstruction, several recent efforts have extended pre-trained diffusion models to tackle scene-level 3D tasks, such as optical flow estimation [69], view synthesis [10, 15, 44, 68, 81, 109], depth estimation [13, 28, 117], and normal estimation [14, 33, 63]. More related to our approach, Matrix3D [49] jointly predicts depth and camera parameters, and WVD [115] introduces a hybrid RGB+point map representation for scene reconstruction. However, these approaches assume *static* 3D environments, whereas we address *dynamic* 4D scene reconstruction, which is a much harder problem due to object motion across time.

More closely related to our approach, concurrent GeometryCrafter [103] introduced a point map VAE with a dual encoder-decoder architecture to improve reconstruction accuracy. However, their point maps are defined in individual camera coordinates, necessitating the use of additional segmentation [30] and tracking models [101] to recover the global point map and estimate camera poses. Aether [82], on the other hand, outputs depth maps and ray maps from a video diffusion model for 4D reconstruction. In contrast, our experiments demonstrate that performance can be significantly enhanced by jointly predicting multiple geometric modalities that capture diverse dynamic ranges, ensuring better temporal coherence and robustness. Importantly, our approach is self-contained and does *not* rely on external models, enhancing its generality and reliability.

3. Method

Our goal is to learn a neural network f_θ that can reconstruct dynamic 3D scenes from monocular videos. Given as input a monocular video $\mathcal{I} = \{I^i\}_{i=1}^N$ consisting of N frames, where each frame is an RGB image $I^i \in \mathbb{R}^{H \times W \times 3}$, the network f_θ returns a representation of its 4D geometry:

$$f_\theta : \{I^i\}_{i=1}^N \mapsto \{(D^i, X^i, r^i)\}_{i=1}^N. \quad (1)$$

The network computes the disparity map $\mathbf{D}^i \in \mathbb{R}^{H \times W \times 1}$, the viewpoint-invariant point map $\mathbf{X}^i \in \mathbb{R}^{H \times W \times 3}$, and the ray map $\mathbf{r}^i \in \mathbb{R}^{H \times W \times 6}$ for each frame \mathbf{I}^i , $i = 1, \dots, N$. As we discuss in Sec. 3.2, these quantities collectively represent the 4D geometry of a scene, including its dynamic structure and time-varying camera extrinsic and intrinsic parameters. No camera parameters are provided as input; these are implicitly estimated by the model as well.

We implement f_θ as a video diffusion model, where θ are the learnable parameters. We discuss the relevant background on video diffusion models in Sec. 3.1. Then, in Sec. 3.2, we describe how we extend the model to predict the three modalities of the 4D geometry. Finally, in Sec. 3.3, we describe how we fuse and align these modalities to obtain a coherent 4D reconstruction at test time.

3.1. Preliminaries: Video Diffusion Model

Our key insight is that by building on pre-trained video diffusion models, our approach can exploit the strong motion and scene geometry priors inherently encoded within these models. Specifically, we build Geo4D on top of DynamiCrafter [102], a ‘‘foundation’’ video diffusion model. DynamiCrafter is a latent diffusion model [67]: it uses a variational autoencoder (VAE) to obtain a more compact video representation and thus reduce computational complexity. During training, a target sequence $\mathcal{X} = \mathbf{x}^{1:N}$ is first encoded into the latent space using the encoder $\mathbf{z}_0^{1:N} = \mathcal{E}(\mathbf{x}^{1:N})$, and then perturbed by $\mathbf{z}_t^{1:N} = \sqrt{\bar{\alpha}_t} \mathbf{z}_0^{1:N} + \sqrt{1 - \bar{\alpha}_t} \epsilon^{1:N}$, where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is Gaussian noise, and $\bar{\alpha}_t$ is the noise level at step t of T noising steps. The denoising network ϵ_θ is then trained to reverse this noising process by optimizing the following objective:

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}^{1:N}, y), t, \epsilon^{1:N} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \|\epsilon_\theta(\mathbf{z}_t^{1:N}, t, y) - \epsilon^{1:N}\|_2^2, \quad (2)$$

where y is the conditional input. Once trained, the model generates a video prompted by y via iteratively denoising from pure noise $\mathbf{z}_T^{1:N}$, and then decoding the denoised latent with a decoder $\hat{\mathcal{X}} = \mathcal{D}(\hat{\mathbf{z}}_0^{1:N})$.

3.2. Multi-modal Geometric 4D Diffusion

We first provide a more precise description of the 4D multi-modal representation output by our model, and then explain how it is encoded in the latent space for generation.

Multi-modal geometric representations. The dynamic 3D structure of a scene is represented by a sequence of point maps $\{\mathbf{X}^i\}_{i=1}^N$, one for each of its N frames. Let (u, v) denote the pixel coordinates in the image plane. Then, the value $X_{uv}^i \in \mathbb{R}^3$ is the 3D coordinate of the scene point that lands at pixel (u, v) in frame \mathbf{I}^i , expressed in the reference frame of camera $i = 1$. Because the reference frame is fixed and independent of the time-varying viewpoint, we call these point maps *viewpoint-invariant*. The

advantages of this representation are convincingly demonstrated by DUS3R [92]. For a static scene, or by knowing which image pixels correspond to the static part of a scene, knowledge of the point maps allows recovery of the intrinsic and extrinsic camera parameters as well as the scene depth. This is done by solving an optimization problem that aligns the dynamic point maps with a pinhole camera model.

As noted in Sec. 1, while point maps $\{\mathbf{X}^i\}_{i=1}^N$ fully encode the 4D geometry of the scene, they are not effective for all parts of the scene. Their dynamic range is limited, and they are not even defined for points at infinity (e.g. sky). Hence, we consider two additional modalities: disparity maps $\{\mathbf{D}^i\}_{i=1}^N$ and camera ray maps $\{\mathbf{r}^i\}_{i=1}^N$, also encouraged by prior evidence [14, 33, 49] that diffusion models can benefit from learning to predict multiple quantities. Disparity maps are not viewpoint-invariant, but have a better dynamic range than point maps (the disparity is zero for points at infinity). Ray maps represent only the camera parameters and are defined for all image pixels, independent of the scene geometry. For the disparity map, D_{uv}^i is the disparity (inverse depth) of the scene point that lands at pixel (u, v) , as seen in frame \mathbf{I}^i . For the ray map, we adopt Plücker coordinates [75, 97, 118], i.e., $\mathbf{r}_{uv} = (\mathbf{d}_{uv}, \mathbf{m}_{uv})$, where $\mathbf{d}_{uv} = \mathbf{R}^\top \mathbf{K}^{-1}(u, v, 1)^\top$ is the ray direction, and $\mathbf{m}_{uv} = -\mathbf{R}^\top \mathbf{t} \times \mathbf{d}_{uv}$, where $(\mathbf{R}, \mathbf{K}, \mathbf{t})$ are the camera’s rotation, calibration, and translation parameters.

Multi-modal latent encoding. The three modalities come in the form of images and can thus be naturally predicted by the video diffusion architecture. However, this requires first mapping them to the latent space, for which we need suitable versions of the encoder \mathcal{E} and decoder \mathcal{D} from Sec. 3.1. Related prior work [14, 28] for depth prediction simply repurposes a pre-trained image encoder-decoder without modification. We found this to work well for disparity and ray maps, but not for point maps. Hence, for the point maps only, we fine-tune the pre-trained decoder \mathcal{D} using the following objective function [100]:

$$\mathcal{L} = - \sum_{uv} \ln \frac{1}{\sqrt{2} \sigma_{uv}} \exp - \frac{\sqrt{2} \ell_1(\mathcal{D}(\mathcal{E}(\mathbf{X}))_{uv}, \mathbf{X}_{uv})}{\sigma_{uv}}, \quad (3)$$

where $\sigma \in \mathbb{R}^{H \times W}$ is the uncertainty of the reconstructed point map, which is also predicted by an additional branch of our VAE decoder. We leave the encoder \mathcal{E} unchanged to modify the latent space as little as possible; instead, we normalize the point maps to the range $[-1, 1]$ to make them more compatible with the pre-trained image encoder.

Video conditioning. The original video diffusion model is conditioned on a single image, but here we need to condition it on the entire input video $\mathcal{I} = \{\mathbf{I}^i\}_{i=1}^N$. To this end, we use a hybrid conditioning mechanism with two streams.

As shown in Fig. 2, in one stream, we extract a global representation of each frame \mathbf{I}^i by passing it to

CLIP [64] followed by a lightweight learnable query transformer [1]. These vectors are incorporated in the transformer via cross-attention layers injected in each U-Net block. In the other stream, we extract *local spatial* features from the VAE encoder and concatenate them channel-wise to the noised latents, encoding the generated 4D modalities $\{(D^i, X^i, r^i)\}_{i=1}^N$.

3.3. Multi-Modal Alignment

As noted, Geo4D predicts several non-independent geometric modalities. Furthermore, processing all frames of a long monocular video simultaneously with a video diffusion model is computationally prohibitive. Therefore, during inference, we use a *temporal sliding window* that segments the video into multiple overlapping clips, with partial overlap to facilitate joining them. The goal of this section is to fuse the resulting multi-modal and multi-window data into a single, coherent reconstruction of the entire video.

Temporal sliding window. Given a video $\mathcal{I} = \{I^i\}_{i=1}^N$ with N frames, we divide it into several video clips $\mathcal{G} = \{g^k\}$, $k \in S$, where each clip g^k contains V frames $\{I^i\}_{i=k}^{k+V-1}$, and the set of starting indices is $S = \{0, s, 2s, \dots, \lfloor \frac{N-V}{s} \rfloor s\} \cup \{N-V\}$. Here, s is the sliding window stride. The final term $\{N-V\}$ ensures that the last clip always includes the final frames of the video.

Alignment objectives. First, given the predicted point maps $X^{i,g}$ for each frame i in each video clip $g \in \mathcal{G}$, we derive corresponding *globally aligned* point maps in world coordinates, as well as the relative camera motion and scale parameters. We denote these quantities with the p subscript to emphasize that they are inferred from the point map predictions. To do so, we extend the pairwise global alignment loss from DUST3R to a *group-wise* one:

$$\mathcal{L}_p(X, \lambda_p^g, P_p^g) = \sum_{g \in \mathcal{G}} \sum_{i \in g} \sum_{uv} \left\| \frac{X_{uv}^i - \lambda_p^g P_p^g X_{uv}^{i,g}}{\sigma_{uv}^{i,g}} \right\|_1, \quad (4)$$

where λ_p^g and $P_p^g = [R_p^g \mid \beta_p^g]$ denote the group-wise scale and transformation matrix that align the group-relative point maps $X^{i,g}$ to the point maps X^i expressed in the global reference frame. $\sigma_{uv}^{i,g}$ denotes the uncertainty of the point map for frame i in group g at pixel (u, v) . We further parameterize each of these point maps as $X_{uv}^i = R_p^{i\top} K_p^{i-1} D_{p,uv}^{i-1}(u, v, 1) + o_p^i$ in terms of each camera's calibration K_p^i , world-to-camera rotation R_p^i , and center o_p^i expressed in the global reference frame, and the disparity map D_p^i . Substituting this expression into the loss function (4) and minimizing it, we can thus recover $K_p^i, R_p^i, o_p^i, D_p^i, \lambda_p^g, P_p^g$ from the predicted point maps.

The steps above infer the disparity maps D_p^i from the point maps, but the model also predicts disparity maps D_d^i directly, where the d subscript denotes disparity prediction.

We introduce the following loss to align them:

$$\mathcal{L}_d(D_p, \lambda_d^g, \beta_d^g) = \sum_{g \in \mathcal{G}} \sum_{i \in g} \left\| D_p^i - \lambda_d^g D_d^{i,g} - \beta_d^g \right\|_1, \quad (5)$$

where λ_d^g and β_d^g are optimized scale and shift parameters.

Finally, the ray maps r also encode camera pose. To align them with the global camera parameters (R_p, K_p, o_p) obtained from the point map, we first solve an optimization problem to extract the camera parameters from the ray map $r^{i,g} = (d^{i,g}, m^{i,g})$ for each group g at frame i . Following Ray Diffusion [114], the camera center $o_c^{i,g}$ is solved by finding the 3D world coordinate closest to the intersection of all rays:

$$o_c^{i,g} = \arg \min_{p \in \mathbb{R}^3} \sum_{u \in H, v \in W} \|p \times d_{uv}^{i,g} - m_{uv}^{i,g}\|^2. \quad (6)$$

The camera extrinsics are solved by optimizing for the matrix H that transforms the predicted per-pixel ray directions $d_{uv}^{i,g}$ to the ray directions u_{uv} of a canonical camera:

$$H^{i,g} = \arg \min_{\|H\|=1} \sum_{u \in H, v \in W} \|H d_{uv}^{i,g} \times u_{uv}\|. \quad (7)$$

Then the world-to-camera rotation matrix $R_c^{i,g}$ and intrinsic matrix $K_c^{i,g}$ can be solved using the RQ-decomposition of $H^{i,g}$. Finally, the camera trajectory alignment loss is:

$$\mathcal{L}_c(R_p, o_p, R_c^g, \beta_c^g, \lambda_c^g) = \sum_{g \in \mathcal{G}} \sum_{i \in g} \left(\left\| R_p^{i\top} R_c^g R_c^{i,g} - I \right\|_f + \left\| \lambda_c^g o_c^{i,g} + \beta_c^g - o_p^i \right\|_2 \right), \quad (8)$$

where $R_c^g, \beta_c^g, \lambda_c^g$ are learnable group-wise rotation matrix, translation vector, and scale, respectively, to align the global camera trajectory (R_p, o_p) and the predicted ones (R_c, o_c) . Following MonST3R [113], we also use a loss to smooth the camera trajectory:

$$\mathcal{L}_s(R_p, o_p) = \sum_{i=1}^N \left(\left\| R_p^{i\top} R_p^{i+1} - I \right\|_f + \left\| o_p^{i+1} - o_p^i \right\|_2 \right). \quad (9)$$

The final optimization objective is the weighted combination of the losses above:

$$\mathcal{L}_{\text{all}} = \alpha_1 \mathcal{L}_p + \alpha_2 \mathcal{L}_d + \alpha_3 \mathcal{L}_c + \alpha_4 \mathcal{L}_s. \quad (10)$$

A note on the invariants. The model predicts point maps, disparity maps, and ray map origins up to scale, as this cannot be uniquely determined from a monocular video. The disparity map is also recovered up to a translation, which discounts the focal length (this is sometimes difficult to estimate due to the dolly zoom effect). Likewise, the ray map origin is recovered up to a shift, necessary to allow normalizing these maps.

Category	Method	Sintel [5]		Bonn [55]		KITTI [17]	
		Abs Rel ↓	$\delta < 1.25 \uparrow$	Abs Rel ↓	$\delta < 1.25 \uparrow$	Abs Rel ↓	$\delta < 1.25 \uparrow$
Single-frame depth	Marigold [28]	0.532	51.5	0.091	93.1	0.149	79.6
	Depth-Anything-V2 [106]	0.367	55.4	0.106	92.1	0.140	80.4
Video depth	NVDS [95]	0.408	48.3	0.167	76.6	0.253	58.8
	ChronoDepth [70]	0.687	48.6	0.100	91.1	0.167	75.9
	DepthCrafter* [22]	<u>0.270</u>	<u>69.7</u>	0.071	97.2	<u>0.104</u>	<u>89.6</u>
Video depth & Camera pose	Robust-CVD [32]	0.703	47.8	—	—	—	—
	CasualSAM [116]	0.387	54.7	0.169	73.7	0.246	62.2
	MonST3R [113]	0.335	58.5	<u>0.063</u>	96.4	<u>0.104</u>	89.5
	Geo4D (Ours)	0.205	73.5	0.059	97.2	0.086	93.7

Table 1. **Video depth estimation** on Sintel [5], Bonn [55] and KITTI [17] datasets. We follow the evaluation protocols established in recent MonST3R [113] for a fair comparison. Notably, results for DepthCrafter* are reported from its latest version (v1.0.1). The **Best** and the **second best** results are highlighted.

4. Experiments

4.1. Experimental Settings

Training datasets. Geo4D is trained exclusively on synthetic datasets, yet demonstrates strong generalization to real-world videos. Specifically, we use five synthetic datasets for training: Spring [50], BEDLAM [2], PointOdyssey [119], TarTanAir [93], and VirtualKitti [6]. See the Supp. Mat Tab. 5 for details.

Training. Our Geo4D is initialized with the weights of DynamiCrafter [102] and trained using AdamW [48] with a learning rate of 1×10^{-5} and a batch size of 32. We use a progressive training strategy to improve convergence and stability. First, we train the model to generate a single geometric modality, *i.e.*, the point maps, at a fixed resolution of 512×320 . Next, we introduce a multi-resolution training scheme to improve generalization and robustness, which includes various resolutions: 512×384 , 512×320 , 576×256 , 640×192 . Finally, we progressively add additional geometric modalities, *i.e.*, the ray and depth maps. Training is conducted on 4 NVIDIA H100 GPUs with a total training time of approximately one week.

Inference. As described in Sec. 3.2, given an N -frame video as input, we first split it into overlapping clips \mathcal{G} , each containing $V = 16$ frames, with a stride of $s = 4$. Each video clip is encoded and fed to the diffusion model to sample multi-modal 4D parameters $(\mathbf{X}^{i,g}, \mathbf{D}^{i,g}, \mathbf{r}^{i,g})$ for the video. For sampling, we use DDIM [77] with 5 steps. Finally, the alignment algorithm in Sec. 3.2 is used to fuse the clips into a globally coherent 4D reconstruction of the entire video.

4.2. Video Depth Estimation

Testing data. Our hypothesis is that, despite being trained on synthetic data, our model can generalize well to out-of-distribution synthetic *and* real data, as it is based on a pre-trained video diffusion model. To test this hypothe-

sis, we evaluate our model on three benchmarks: Sintel [5] is a synthetic dataset that provides accurate depth annotations, covering diverse scenes with complex camera motion. KITTI [17] is a large driving dataset collected using stereo cameras and LiDAR sensors. Bonn [55] focuses on dynamic indoor scenes. To ensure fair comparisons, we follow the evaluation protocol used by MonST3R [113], where depth sequences are uniformly sampled from the datasets, extracting 50–110 frames per sequence for evaluation.

Metrics. Following the standard affine-invariant depth evaluation protocol [65], we align the predicted video depth with the ground-truth depth before computing metrics. However, unlike single-image depth estimation [28, 105, 106], where depth alignment is performed per frame, we enforce *global scale consistency* by applying a single scale and shift across the entire video sequence. For quantitative evaluation, we adopt two widely used depth metrics: absolute relative error (Abs Rel) and the percentage of inlier points (with a threshold value of $\delta < 1.25$).

Baselines. We compare Geo4D to state-of-the-art single-frame depth estimation methods (Marigold [28] and Depth-Anything-V2 [106]), video depth prediction (NVDS [95], ChronoDepth [70], and DepthCrafter [22]), and joint video depth and camera pose prediction (Robust-CVD [32], CausalSAM [116], and MonST3R [113]).

Results. As shown in Table 1, all versions of Geo4D outperform state-of-the-art methods by a large margin. This includes DepthCrafter [22] and MonST3R [113], the most recent video depth diffusion model and the dynamic extension of DUST3R to dynamic scenes, respectively. Notably, while both Geo4D and DepthCrafter are based on the same video diffusion model (DynamiCrafter), our model outperforms DepthCrafter in Abs Rel by 24.0% on Sintel and 17.3% on KITTI, despite solving a more general problem. Qualitatively, Fig. 3 shows that Geo4D achieves more consistent results, especially for fast-moving objects.

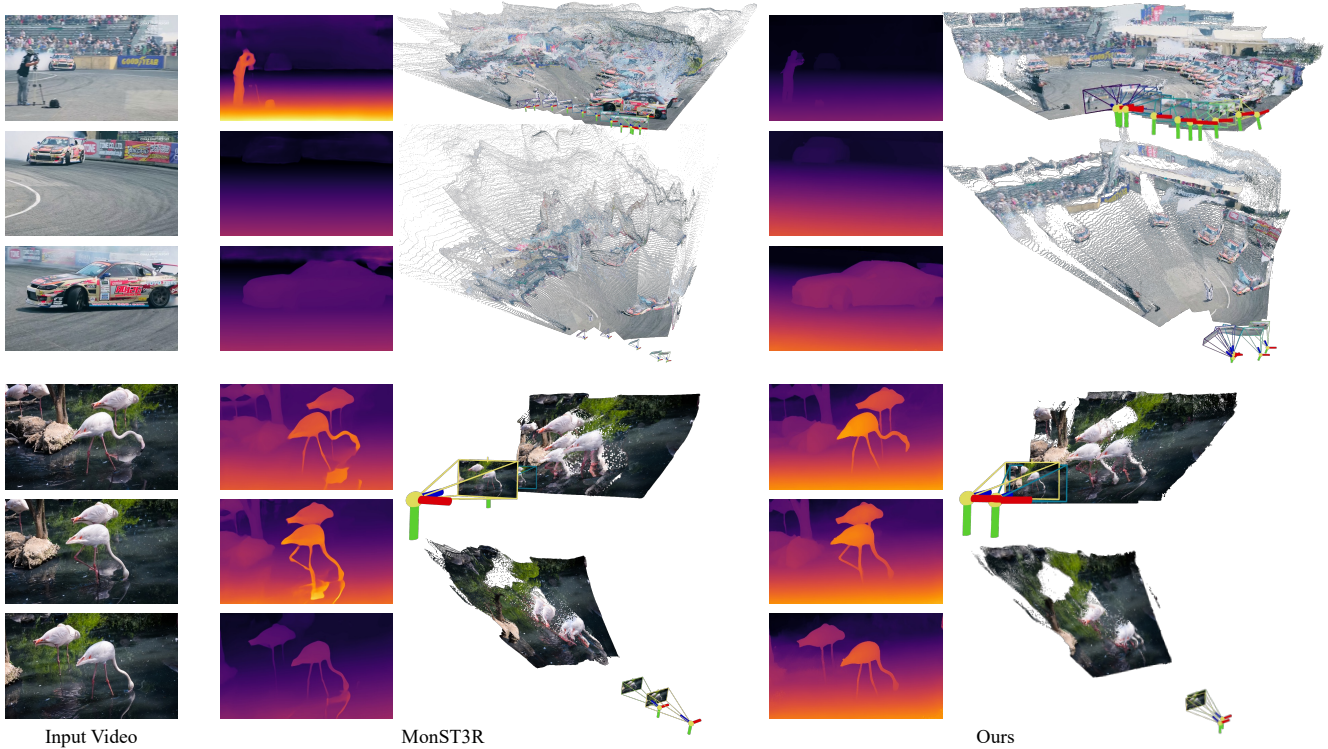


Figure 3. **Qualitative results** comparing Geo4D with MonST3R [113]. Attributed to our group-wise inference manner and prior geometry knowledge from pretrained video diffusion, our model successfully produces consistent 4D geometry under fast motion (first row) and deceptive reflection in the water (second row).

4.3. Camera Pose Estimation

Setup. We evaluate the performance of Geo4D on both the synthetic Sintel [5] dataset and the realistic TUM-dynamics [78] dataset. We follow the same evaluation protocol as in MonST3R [113]. Specifically, on Sintel, we select 14 dynamic sequences, and for TUM-dynamics, we sample the first 90 frames of each sequence with a temporal stride of 3. After aligning the predicted camera trajectory with the ground truth using the Umeyama algorithm, we calculate three commonly used metrics: Absolute Translation Error (ATE), Relative Translation Error (RPE-T), and Relative Rotation Error (RPE-R). We compare our method with other state-of-the-art discriminative methods, which jointly predict camera pose and depth, including Robust-CVD [32], CausalSAM [116], and MonST3R [113].

Results. To the best of our knowledge, Geo4D is the first method that uses a generative model to estimate camera parameters in a dynamic scene. As shown in Tab. 2, compared to existing non-generative alternatives, we achieve much better camera rotation prediction (RPE-R) and comparable camera translation prediction (ATE and RPE-T).

4.4. Qualitative Comparison

4D reconstruction. We compare Geo4D with the state-of-the-art MonST3R method on the DAVIS [23] dataset. Up-

Method	Sintel			TUM-dynamics		
	ATE ↓	RPE-T ↓	RPE-R ↓	ATE ↓	RPE-T ↓	RPE-R ↓
Robust-CVD [32]	0.360	0.154	3.443	0.153	0.026	3.528
CausalSAM [116]	0.141	0.035	0.615	0.071	0.010	1.712
MonST3R [113]	0.108	0.042	0.732	0.063	0.009	1.217
Geo4D (Ours)	0.185	0.063	0.547	0.073	0.020	0.635

Table 2. **Quantitative evaluation for camera pose estimation.** We achieve comparable camera pose estimation performance with other discriminative SOTA methods.

grading from pairwise alignment as in MonST3R to our group-wise alignment improves temporal consistency, leading to a more stable and globally coherent 4D reconstruction of point maps and camera trajectory, particularly in highly dynamic scenes. As shown in the top row of Fig. 3, Geo4D successfully tracks the racing car in 4D, whereas MonST3R struggles due to the rapid motion between pairs of images. Furthermore, likely due to the strong prior captured by the pre-trained video generative model, Geo4D correctly reconstructs the reflection of the flamingo in the water (second row in Fig. 3), whereas MonST3R misinterprets the reflection as a foreground object, resulting in incorrect depth.

Video depth prediction. We compare Geo4D with state-of-the-art video depth predictors MonST3R [113] and DepthCrafter [22] on the Sintel [5] dataset. Qualitatively, Geo4D produces more detailed geometry, for instance for




Figure 4. **Qualitative video depth results** comparing Geo4D with MonST3R [113] and DepthCrafter [22]. Owing to our proposed multi-modal training and alignment, as well as the prior knowledge from diffusion, our method can infer a more detailed structure (first row) and a more accurate spatial arrangement from video (second row).

Training			Inference			Video Depth		Camera Pose		
Point Map	Disparity Map	Ray Map	Point Map	Disparity Map	Ray Map	Abs Rel ↓	$\delta < 1.25$ ↑	ATE ↓	RPE trans ↓	RPE rot ↓
✓	—	—	✓	—	—	0.232	71.3	0.335	0.076	0.731
✓	✓	✓	✓	—	—	0.223	72.5	0.237	0.070	0.566
✓	✓	✓	—	✓	—	0.211	73.4	—	—	—
✓	✓	✓	—	—	✓	—	—	0.268	0.192	1.476
✓	✓	✓	✓	✓	✓	0.205	73.5	0.185	0.063	0.547

Table 3. **Ablation study for the different modalities of the geometric representation** on the Sintel [5] dataset. We demonstrate the effectiveness of our key design choices that both leverage multi-modality as additional training supervision signal and postprocess through our proposed multi-modal alignment algorithm will improve the overall performance.

Stride	s / frame	Video Depth		Camera Pose		
		Abs Rel ↓	$\delta < 1.25$ ↑	ATE ↓	RPE trans ↓	RPE rot ↓
15	0.92	0.213	72.4	0.210	0.092	0.574
8	1.24	0.212	72.8	0.222	0.074	0.524
4	1.89	0.205	73.5	0.185	0.063	0.547
2	3.26	0.204	72.9	0.181	0.058	0.518

Table 4. **Ablation study for the temporal sliding window stride** on the Sintel [5] dataset. There is a trade-off between performance and inference speed.

the rope on the stick in the first row of Fig. 4, and a better spatial arrangement between different dynamic objects, as shown in the second row of Fig. 4.

4.5. Ablation Study

We ablate our key design choices and the effect of different modalities on the Sintel dataset.

We study the effect of multi-modality in Tab. 3. The three modalities—point map, disparity map, and ray map—can be used either at training or inference time, or both. The first two rows show that the diffusion model trained with point maps as a single modality performs worse in both video depth and camera pose estimation than the diffusion model trained with all three modalities. Therefore, the other two modalities, even if they can be seen as redundant, serve as additional supervisory signals during training, which improves the generalization ability of the diffusion model.

We then investigate the effectiveness of our multi-modal alignment algorithm. Compared with the second to the fourth row in Tab. 3, which leverage only a single modality during inference, multi-modal alignment optimization (last row) achieves the best performance, showing the benefits of fusing the multiple modalities at inference time.

We ablate the sliding window stride in Tab. 4. Results improve with a shorter stride, in part because this means that more windows and estimates are averaged, reducing the variance of the predictions by the denoising diffusion model, which is stochastic. We choose stride $s = 4$ for our main results to balance runtime and performance. Note that MonST3R [113] requires 2.41 seconds to process one frame under the same setting, so our method is 1.27 times faster than MonST3R [113].

5. Discussion and Conclusion

We have introduced Geo4D, a novel approach that adapts a video generator for dynamic 4D reconstruction. By building on a pre-trained video generator, Geo4D achieves excellent generalization to real data despite being trained only on synthetic 4D data. We have also demonstrated the benefits of predicting multiple modalities and fusing them at test time via optimization. Our model outperforms state-of-the-art methods on video depth and camera rotation prediction, particularly in challenging dynamic scenes.

Despite these successes, our approach has limitations. One is that the point map encoder-decoder is still not entirely accurate, which in turn is a bottleneck for the overall reconstruction quality.

Our approach also opens a path to integrating 4D geometry into video foundation models, *e.g.*, to generate 3D animations from text, or to provide a more actionable signal when the video model is used as a proxy for a world model.

Acknowledgments. The authors of this work were supported by Clarendon Scholarship, ERC 101001212-UNION, and EPSRC EP/Z001811/1 SYN3D.

References

- [1] Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, Jenia Jitsev, Simon Kornblith, Pang Wei Koh, Gabriel Ilharco, Mitchell Wortsman, and Ludwig Schmidt. Openflamingo: An open-source framework for training large autoregressive vision-language models, 2023. 5
- [2] Michael J. Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. Bedlam: A synthetic dataset of bodies exhibiting detailed lifelike animated motion, 2023. 6, 15
- [3] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelvitich, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 3
- [4] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 22563–22575, 2023. 3
- [5] Daniel J. Butler, Jonas Wulff, Garrett B. Stanley, and Michael J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision (ECCV)*, 2012. 6, 7, 8, 15, 16
- [6] Johann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2, 2020. 6, 15
- [7] David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19457–19467, 2024. 2
- [8] Xingyu Chen, Yue Chen, Yuliang Xiu, Andreas Geiger, and Anpei Chen. Easi3r: Estimating disentangled motion from dust3r without training. *arXiv preprint arXiv:2503.24391*, 2025. 3
- [9] Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. MVSplat: efficient 3d gaussian splatting from sparse multi-view images. *arXiv*, 2403.14627, 2024. 2
- [10] Yuedong Chen, Chuanxia Zheng, Haofei Xu, Bohan Zhuang, Andrea Vedaldi, Tat-Jen Cham, and Jianfei Cai. Mvsplat360: Feed-forward 360 scene synthesis from sparse views. In *Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [11] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision (ECCV)*, pages 628–644. Springer, 2016. 2
- [12] Wen-Hsuan Chu, Lei Ke, and Katerina Fragkiadaki. Dreamscene4d: Dynamic multi-object scene generation from monocular videos. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2
- [13] Yiquan Duan, Xianda Guo, and Zheng Zhu. Diffusion-depth: Diffusion denoising approach for monocular depth estimation. In *European Conference on Computer Vision (ECCV)*, pages 432–449. Springer, 2024. 3
- [14] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. In *European Conference on Computer Vision (ECCV)*, pages 241–258. Springer, 2024. 3, 4
- [15] Ruiqi Gao, Aleksander Holynski, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul Srinivasan, Jonathan T Barron, and Ben Poole. Cat3d: Create anything in 3d with multi-view diffusion models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [16] Songwei Ge, Seungjun Nah, Guilin Liu, Tyler Poon, Andrew Tao, Bryan Catanzaro, David Jacobs, Jia-Bin Huang, Ming-Yu Liu, and Yogesh Balaji. Preserve your own correlation: A noise prior for video diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22930–22941, 2023. 3
- [17] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research (IJRR)*, 2013. 6
- [18] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *Proceedings of the IEEE/CVF international conference on computer vision (CVPR)*, pages 9785–9795, 2019. 2
- [19] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. In *International Conference on Learning Representations (ICLR)*, 2024. 3
- [20] Junlin Han, Filippos Kokkinos, and Philip Torr. Vfusion3d: Learning scalable 3d generative models from video diffusion models. In *European Conference on Computer Vision (ECCV)*, pages 333–350. Springer, 2024. 3
- [21] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Neural Information Processing Systems (NeurIPS)*, 35:8633–8646, 2022. 3
- [22] Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiadong Cun, Yong Zhang, Long Quan, and Ying Shan. Depthrafter: Generating consistent long depth sequences for open-world videos. In *CVPR*, 2025. 2, 6, 7, 8, 15
- [23] Jia-Bin Huang, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. Temporally coherent completion of dynamic video. In *ACM*, 2016. 2, 7
- [24] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. Volumedeform: Real-time volumetric non-rigid reconstruction. In *European conference on computer vision (ECCV)*, pages 362–379. Springer, 2016. 2
- [25] Tomas Jakab, Ruining Li, Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Farm3d: Learning articulated 3d animals by distilling 2d diffusion. In *2024 International*

- Conference on 3D Vision (3DV)*, pages 852–861. IEEE, 2024. 3
- [26] Yanqin Jiang, Chaohui Yu, Chenjie Cao, Fan Wang, Weiming Hu, and Jin Gao. Animate3d: Animating any 3d model with multi-view video diffusion. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [27] Zeren Jiang, Chen Guo, Manuel Kaufmann, Tianjian Jiang, Julien Valentin, Otmar Hilliges, and Jie Song. Multiply: Reconstruction of multiple people from monocular video in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [28] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9492–9502, 2024. 2, 3, 4, 6
- [29] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2, 3
- [30] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything, 2023. 3
- [31] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. 3
- [32] Johannes Kopf, Xuejian Rong, and Jia-Bin Huang. Robust consistent video depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1611–1621, 2021. 6, 7
- [33] Akshay Krishnan, Xinchun Yan, Vincent Casser, and Abhijit Kundu. Orchid: Image latent diffusion for joint appearance and geometry generation. *arXiv preprint arXiv:2501.13087*, 2025. 3, 4
- [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
- [35] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, pages 71–91. Springer, 2024. 2
- [36] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3D: Fast text-to-3D with sparse-view generation and large reconstruction model. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024. 3
- [37] Xuanyi Li, Daquan Zhou, Chenxu Zhang, Shaodong Wei, Qibin Hou, and Ming-Ming Cheng. Sora generates videos with stunning geometrical consistency. *arXiv*, 2402.17403, 2024. 2
- [38] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6498–6508, 2021. 2
- [39] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snavely. Dynibar: Neural dynamic image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4273–4284, 2023. 2
- [40] Zhengqi Li, Richard Tucker, Forrester Cole, Qianqian Wang, Linyi Jin, Vickie Ye, Angjoo Kanazawa, Aleksander Holynski, and Noah Snavely. Megasam: Accurate, fast and robust structure and motion from casual dynamic videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [41] Chen-Hsuan Lin, Chen Kong, and Simon Lucey. Learning efficient point cloud generation for dense 3d object reconstruction. In *proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2018. 2
- [42] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 300–309, 2023. 3
- [43] Youtian Lin, Zuozhuo Dai, Siyu Zhu, and Yao Yao. Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In *CVPR*, 2024. 2
- [44] Fangfu Liu, Wenqiang Sun, Hanyang Wang, Yikai Wang, Haowen Sun, Junliang Ye, Jun Zhang, and Yueqi Duan. Reconx: Reconstruct any scene from sparse views with video diffusion model, 2024. 3
- [45] 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 *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*, pages 9298–9309, 2023. 3
- [46] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024.
- [47] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 9970–9980, 2024. 3
- [48] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 6
- [49] Yuanxun Lu, Jingyang Zhang, Tian Fang, Jean-Daniel Nahmias, Yanghai Tsin, Long Quan, Xun Cao, Yao Yao, and Shiwei Li. Matrix3d: Large photogrammetry model all-in-one. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3, 4

- [50] Lukas Mehl, Jenny Schmalfluss, Azin Jahedi, Yaroslava Nalivayko, and Andrés Bruhn. Spring: A high-resolution high-detail dataset and benchmark for scene flow, optical flow and stereo. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 6, 15
- [51] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 8446–8455, 2023. 3
- [52] B Mildenhall, PP Srinivasan, M Tancik, JT Barron, R Ramamoorthi, and R Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision (ECCV)*, 2020. 2
- [53] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 343–352, 2015. 2
- [54] NVIDIA, Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, Daniel Dworakowski, Jiaojiao Fan, Michele Fenzi, Francesco Ferroni, Sanja Fidler, Dieter Fox, Songwei Ge, Yunhao Ge, Jinwei Gu, Siddharth Gururani, Ethan He, Jiahui Huang, Jacob Huffman, Pooya Jannaty, Jingyi Jin, Seung Wook Kim, Gergely Klár, Grace Lam, Shiyi Lan, Laura Leal-Taixe, Anqi Li, Zhaoshuo Li, Chen-Hsuan Lin, Tsung-Yi Lin, Huan Ling, Ming-Yu Liu, Xian Liu, Alice Luo, Qianli Ma, Hanzi Mao, Kaichun Mo, Arsalan Mousavian, Seungjun Nah, Sriharsha Niverty, David Page, Despoina Paschalidou, Zeeshan Patel, Lindsey Pavao, Morteza Ramezani, Fitsum Reda, Xiaowei Ren, Vasanth Rao Naik Sabavat, Ed Schmerling, Stella Shi, Bartosz Stefaniak, Shitao Tang, Lyne Tchampt, Przemek Tredak, Wei-Cheng Tseng, Jibin Varghese, Hao Wang, Haoxiang Wang, Heng Wang, Ting-Chun Wang, Fangyin Wei, Xinyue Wei, Jay Zhangjie Wu, Jiashu Xu, Wei Yang, Lin Yen-Chen, Xiaohui Zeng, Yu Zeng, Jing Zhang, Qinsheng Zhang, Yuxuan Zhang, Qingqing Zhao, and Artur Zolkowski. Cosmos world foundation model platform for physical ai. *arXiv*, 2501.03575, 2025. 2
- [55] Emanuele Palazzolo, Jens Behley, Philipp Lottes, Philippe Giguère, and C. Stachniss. Refusion: 3d reconstruction in dynamic environments for rgb-d cameras exploiting residuals. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7855–7862, 2019. 6
- [56] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019. 2
- [57] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5865–5874, 2021. 2
- [58] 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
- [59] Jack Parker-Holder, Philip Ball, Jake Bruce, Vibhavari Dasagi, Kristian Holsheimer, Christos Kaplanis, Alexandre Moufaret, Guy Scully, Jeremy Shar, Jimmy Shi, Stephen Spencer, Jessica Yung, Michael Dennis, Sultan Kenjeyev, Shangbang Long, Vlad Mnih, Harris Chan, Maxime Gazeau, Bonnie Li, Fabio Pardo, Luyu Wang, Lei Zhang, Frederic Besse, Tim Harley, Anna Mitenkova, Jane Wang, Jeff Clune, Demis Hassabis, Raia Hadsell, Adrian Bolton, Satinder Singh, and Tim Rocktäschel. Genie 2: A large-scale foundation world model, 2024. 2
- [60] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European conference on computer vision (ECCV)*, pages 523–540. Springer, 2020. 2
- [61] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023. 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 (CVPR)*, pages 10318–10327, 2021. 2
- [63] Lingteng Qiu, Guanying Chen, Xiaodong Gu, Qi Zuo, Mutian Xu, Yushuang Wu, Weihao Yuan, Zilong Dong, Liefeng Bo, and Xiaoguang Han. Richdreamer: A generalizable normal-depth diffusion model for detail richness in text-to-3d. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 9914–9925, 2024. 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 (ICML)*, pages 8748–8763. PMLR, 2021. 5
- [65] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44:1623–1637, 2019. 6
- [66] Jiawei Ren, Kevin Xie, Ashkan Mirzaei, Hanxue Liang, Xiaohui Zeng, Karsten Kreis, Ziwei Liu, Antonio Torralba, Sanja Fidler, Seung Wook Kim, and Huan Ling. L4gm: Large 4d gaussian reconstruction model. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [67] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 10684–10695, 2022. 4

- [68] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, et al. Zeronvs: Zero-shot 360-degree view synthesis from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9420–9429, 2024. 3
- [69] Saurabh Saxena, Charles Herrmann, Junhwa Hur, Abhishek Kar, Mohammad Norouzi, Deqing Sun, and David J Fleet. The surprising effectiveness of diffusion models for optical flow and monocular depth estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 36:39443–39469, 2023. 3
- [70] Jiahao Shao, Yuanbo Yang, Hongyu Zhou, Youmin Zhang, Yujun Shen, Matteo Poggi, and Yiyi Liao. Learning temporally consistent video depth from video diffusion priors. *arXiv*, 2406.01493, 2024. 6
- [71] Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. MVDream: Multi-view diffusion for 3d generation. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024. 3
- [72] Yawar Siddiqui, Antonio Alliegro, Alexey Artemov, Tatiana Tommasi, Daniele Sirigatti, Vladislav Rosov, Angela Dai, and Matthias Nießner. Meshgpt: Generating triangle meshes with decoder-only transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19615–19625, 2024. 2
- [73] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, Devi Parikh, Sonal Gupta, and Yaniv Taigman. Make-a-video: Text-to-video generation without text-video data. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023. 3
- [74] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2437–2446, 2019. 2
- [75] Vincent Sitzmann, Semon Rezhikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:19313–19325, 2021. 2, 4
- [76] Brandon Smart, Chuanxia Zheng, Iro Laina, and Victor Adrian Prisacariu. Splatt3r: Zero-shot gaussian splatting from uncalibrated image pairs. *arXiv preprint arXiv:2408.13912*, 2024. 2
- [77] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022. 6
- [78] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of rgb-d slam systems. *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 573–580, 2012. 7
- [79] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Splatter Image: Ultra-fast single-view 3D reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [80] Stanislaw Szymanowicz, Eldar Insafutdinov, Chuanxia Zheng, Dylan Campbell, João F. Henriques, Christian Rupprecht, and Andrea Vedaldi. Flash3D: Feed-forward generalisable 3D scene reconstruction from a single image. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2025. 2
- [81] Stanislaw Szymanowicz, Jason Y Zhang, Pratul Srinivasan, Ruiqi Gao, Arthur Brussee, Aleksander Holynski, Ricardo Martin-Brualla, Jonathan T Barron, and Philipp Henzler. Bolt3d: Generating 3d scenes in seconds. *arXiv preprint arXiv:2503.14445*, 2025. 3
- [82] Aether Team, Haoyi Zhu, Yifan Wang, Jianjun Zhou, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Chunhua Shen, Jiangmiao Pang, and Tong He. Aether: Geometric-aware unified world modeling. *arXiv preprint arXiv:2503.18945*, 2025. 3
- [83] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 2626–2634, 2017. 2
- [84] S. Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(4):376–380, 1991. 15
- [85] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. In *European Conference on Computer Vision (ECCV)*, pages 439–457. Springer, 2024. 3
- [86] Hengyi Wang and Lourdes Agapito. 3d reconstruction with spatial memory. In *International Conference on 3D Vision (3DV)*, 2024. 2
- [87] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 12619–12629, 2023. 3
- [88] Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571*, 2023. 3
- [89] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. VGGT: Visual geometry grounded network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 1, 2
- [90] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European conference on computer vision (ECCV)*, pages 52–67, 2018. 2
- [91] Qianqian Wang, Vickie Ye, Hang Gao, Weijia Zeng, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of

- motion: 4d reconstruction from a single video. In *arXiv preprint arXiv:2407.13764*, 2024. 2
- [92] 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 (CVPR)*, pages 20697–20709, 2024. 1, 2, 3, 4
- [93] Wenshan Wang, DeLong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian A. Scherer. Tartanair: A dataset to push the limits of visual slam. *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4909–4916, 2020. 6, 15
- [94] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiuniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Videocomposer: Compositional video synthesis with motion controllability. *Advances in Neural Information Processing Systems (NeurIPS)*, 36:7594–7611, 2023. 3
- [95] Yiran Wang, Min Shi, Jiaqi Li, Zihao Huang, Zhiguo Cao, Jianming Zhang, Ke Xian, and Guosheng Lin. Neural video depth stabilizer. In *ICCV*, 2023. 6
- [96] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in Neural Information Processing Systems (NeurIPS)*, 36, 2024. 3
- [97] Daniel Watson, William Chan, Ricardo Martin Brullalla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023. 4
- [98] Philippe Weinzaepfel, Vincent Leroy, Thomas Lucas, Romain BRÉGIER, Yohann Cabon, Vaibhav ARORA, Leonid Antsfeld, Boris Chidlovskii, Gabriela Csurka, and Jerome Revaud. CroCo: self-supervised pre-training for 3D vision tasks by cross-view completion. In *Proc. NeurIPS*, 2022. 2
- [99] 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. In *CVPR*, 2024. 2
- [100] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 1–10, 2020. 4
- [101] Yuxi Xiao, Qianqian Wang, Shangzhan Zhang, Nan Xue, Sida Peng, Yujun Shen, and Xiaowei Zhou. Spatialtracker: Tracking any 2d pixels in 3d space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [102] Jinbo Xing, Menghan Xia, Yong Zhang, Haoxin Chen, Wangbo Yu, Hanyuan Liu, Gongye Liu, Xintao Wang, Ying Shan, and Tien-Tsin Wong. Dynamicrafter: Animating open-domain images with video diffusion priors. In *European Conference on Computer Vision (ECCV)*, pages 399–417. Springer, 2024. 1, 3, 4, 6
- [103] Tian-Xing Xu, Xiangjun Gao, Wenbo Hu, Xiaoyu Li, Song-Hai Zhang, and Ying Shan. Geometryrafter: Consistent geometry estimation for open-world videos with diffusion priors, 2025. 3
- [104] Jianing Yang, Alexander Sax, Kevin J. Liang, Mikael Henaff, Hao Tang, Ang Cao, Joyce Chai, Franziska Meier, and Matt Feiszli. Fast3r: Towards 3d reconstruction of 1000+ images in one forward pass. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2
- [105] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proc. CVPR*, 2024. 6
- [106] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 6
- [107] 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 (CVPR)*, pages 20331–20341, 2024. 2
- [108] David Yifan Yao, Albert J. Zhai, and Shenlong Wang. Uni4d: Unifying visual foundation models for 4d modeling from a single video, 2025. 3
- [109] Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. Viewrafter: Taming video diffusion models for high-fidelity novel view synthesis. *arXiv preprint arXiv:2409.02048*, 2024. 3
- [110] Xumin Yu, Yongming Rao, Ziyi Wang, Zuyan Liu, Jiwen Lu, and Jie Zhou. PointR: Diverse point cloud completion with geometry-aware transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12498–12507, 2021. 2
- [111] Yuheng Yuan, Qihong Shen, Xingyi Yang, and Xinchao Wang. 1000+ fps 4d gaussian splatting for dynamic scene rendering, 2025. 2
- [112] David Junhao Zhang, Jay Zhangjie Wu, Jia-Wei Liu, Rui Zhao, Lingmin Ran, Yuchao Gu, Difei Gao, and Mike Zheng Shou. Show-1: Marrying pixel and latent diffusion models for text-to-video generation. *International Journal of Computer Vision (IJCV)*, pages 1–15, 2024. 3
- [113] 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. In *International Conference on Learning Representations (ICLR)*, 2025. 1, 2, 3, 5, 6, 7, 8
- [114] Jason Y Zhang, Amy Lin, Moneish Kumar, Tzu-Hsuan Yang, Deva Ramanan, and Shubham Tulsiani. Cameras as rays: Pose estimation via ray diffusion. In *International Conference on Learning Representations (ICLR)*, 2024. 5
- [115] Qihang Zhang, Shuangfei Zhai, Miguel Angel Bautista, Kevin Miao, Alexander Toshev, Joshua Susskind, and Jiantao Gu. World-consistent video diffusion with explicit 3d

- modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. [3](#)
- [116] Zhoutong Zhang, Forrester Cole, Zhengqi Li, Michael Rubinstein, Noah Snavely, and William T. Freeman. Structure and motion from casual videos. In *European Conference on Computer Vision (ECCV)*, 2022. [6](#), [7](#)
- [117] Wenliang Zhao, Yongming Rao, Zuyan Liu, Benlin Liu, Jie Zhou, and Jiwen Lu. Unleashing text-to-image diffusion models for visual perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5729–5739, 2023. [3](#)
- [118] Chuanxia Zheng and Andrea Vedaldi. Free3d: Consistent novel view synthesis without 3d representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9720–9731, 2024. [3](#), [4](#)
- [119] Yang Zheng, Adam W Harley, Bokui Shen, Gordon Wetstein, and Leonidas J Guibas. Pointodyssey: A large-scale synthetic dataset for long-term point tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19855–19865, 2023. [6](#), [15](#)

Geo4D: Leveraging Video Generators for Geometric 4D Scene Reconstruction

Supplementary Material

In this **supplementary material**, we provide additional information to supplement our main submission. The **code** is available here for research purposes: github.com/jzr99/Geo4D

6. Implementation Details

6.1. Training Dataset

As shown in Tab. 5, we use five synthetic datasets for training: Spring [50], BEDLAM [2], PointOdyssey [119], TartanAir [93], and VirtualKitti [6]. Although all datasets are synthetic, we found that some depth pixels are missing in PointOdyssey [119]. To address this, we apply max pooling to inpaint the missing pixels. During training, we sample each dataset according to the ratios in Tab. 5. For each sample, we select 16 frames from the sequence, with the sampling stride randomly chosen from $\{1, 2, 3\}$ to allow our diffusion model to adapt to input videos with various frame rates.

6.2. Optimization Details

The overall optimization process is outlined in Algorithm 1. We first predict all three modality maps using our diffusion model for each video clip g . The predicted point maps are then roughly aligned based on the overlapping frames using the Umeyama algorithm [84]. The camera intrinsic \mathbf{K}^k is initialized by minimizing the projection error of the point map \mathbf{X}^{k,g^k} in its reference (first) frame k within each window group g^k . The camera extrinsics are then initialized using the RANSAC PnP algorithm. In the first stage of optimization, the point maps are roughly disentangled into camera pose and depth map. The disparity map is then aligned with the global depth inferred from point maps by solving Eq. (5) from the main paper to obtain the scale and shift parameters. The camera parameters extracted from the predicted ray map are aligned with the global camera trajectory based on the reference (first) frame of each video clip g via Eq. (8) from the main paper. After initializing all the alignment learnable parameters, including rotation \mathbf{R}_*^g , scale λ_*^g , and shift β_*^g across different modalities, where $* \in \{p, d, c\}$, we jointly optimize all the learnable parameters by Eq. (10). Specifically, we set the weights for each loss term in Eq. (10) as $\alpha_1 = 1, \alpha_2 = 2, \alpha_3 = 0.005, \alpha_4 = 0.015$ to roughly equalize the scale of the different losses.

Algorithm 1 Multi-Modal Alignment Optimization

```
1:  $\mathbf{X}^{i,g}, \mathbf{D}^{i,g}, \mathbf{r}^{i,g} \leftarrow$  Predicted by our diffusion model
2:  $\mathbf{D}_p^i, \lambda_p^g, \mathbf{R}_p^g, \beta_p^g \leftarrow$  Initialized by Umeyama algorithm
3:  $\mathbf{K}_p^k \leftarrow$  Optimized from  $\mathbf{X}^{k,g^k}$ 
4:  $\mathbf{R}_p^i, \mathbf{o}_p^i \leftarrow$  Initialized by Ransac PnP from pointmaps  $\mathbf{X}^i$ 
5:  $\mathbf{R}_c^{i,g}, \mathbf{o}_c^{i,g} \leftarrow$  Initialized by Eqs. (6) and (7) from raymaps  $\mathbf{r}^{i,g}$ 
6: repeat
7:   if Iteration = Align start iteration then
8:      $\lambda_d^g, \beta_d^g \leftarrow \arg \min \mathcal{L}_d$  ( Eq. (5))
9:      $\mathbf{R}_c^g, \lambda_c^g, \beta_c^g \leftarrow \arg \min \mathcal{L}_c$  ( Eq. (8))
10:  else if Iteration < Align start iteration then
11:     $\mathbf{D}_p^i, \mathbf{K}_p^i, \mathbf{R}_p^i, \mathbf{o}_p^i, \lambda_p^g, \mathbf{R}_p^g, \beta_p^g \leftarrow \arg \min \mathcal{L}_p + \mathcal{L}_s$ 
12:  else
13:     $\mathbf{D}_p^i, \mathbf{K}_p^i, \mathbf{R}_p^i, \mathbf{o}_p^i, \lambda_*^g, \mathbf{R}_*^g, \beta_*^g \leftarrow \arg \min \mathcal{L}_{all}$ 
14:  end if
15: until max loop reached
```

Dataset	Scene type	#Frames	#Sequences	Ratio
PointOdyssey [119]	Indoors/Outdoors	200K	131	16.7%
TartanAir [93]	Indoors/Outdoors	1000K	163	16.7%
Spring [50]	Outdoors	6K	37	16.7%
VirtualKITTI [6]	Driving	43K	320	16.7%
BEDLAM [2]	Indoors/Outdoors	380K	10K	33.3%

Table 5. **Details of training datasets.** Our method only uses synthetic datasets for training.

Steps	Video Depth		Camera Pose		
	Abs Rel ↓	$\delta < 1.25 \uparrow$	ATE ↓	RPE trans ↓	RPE rot ↓
1	0.221	70.7	0.234	0.072	0.753
5	0.205	73.5	0.185	0.063	0.547
10	0.207	73.2	0.212	0.071	0.508
25	0.220	72.2	0.211	0.074	0.564

Table 6. **Ablation study for the DDIM sampling steps.** on the Sintel [5] dataset.

7. Additional Analysis

7.1. Ablating the Number of Denoising Steps

We study the influence of the number of denoising steps during inference. As shown in Tab. 6, the model achieves optimal performance after around 5 steps. Compared to the video generation task, where a larger number of denoising steps usually produces a more detailed generated video, 4D reconstruction is a more deterministic task, which requires fewer steps. Similar phenomena are also observed in [22], which uses a video generator for video depth estimation.

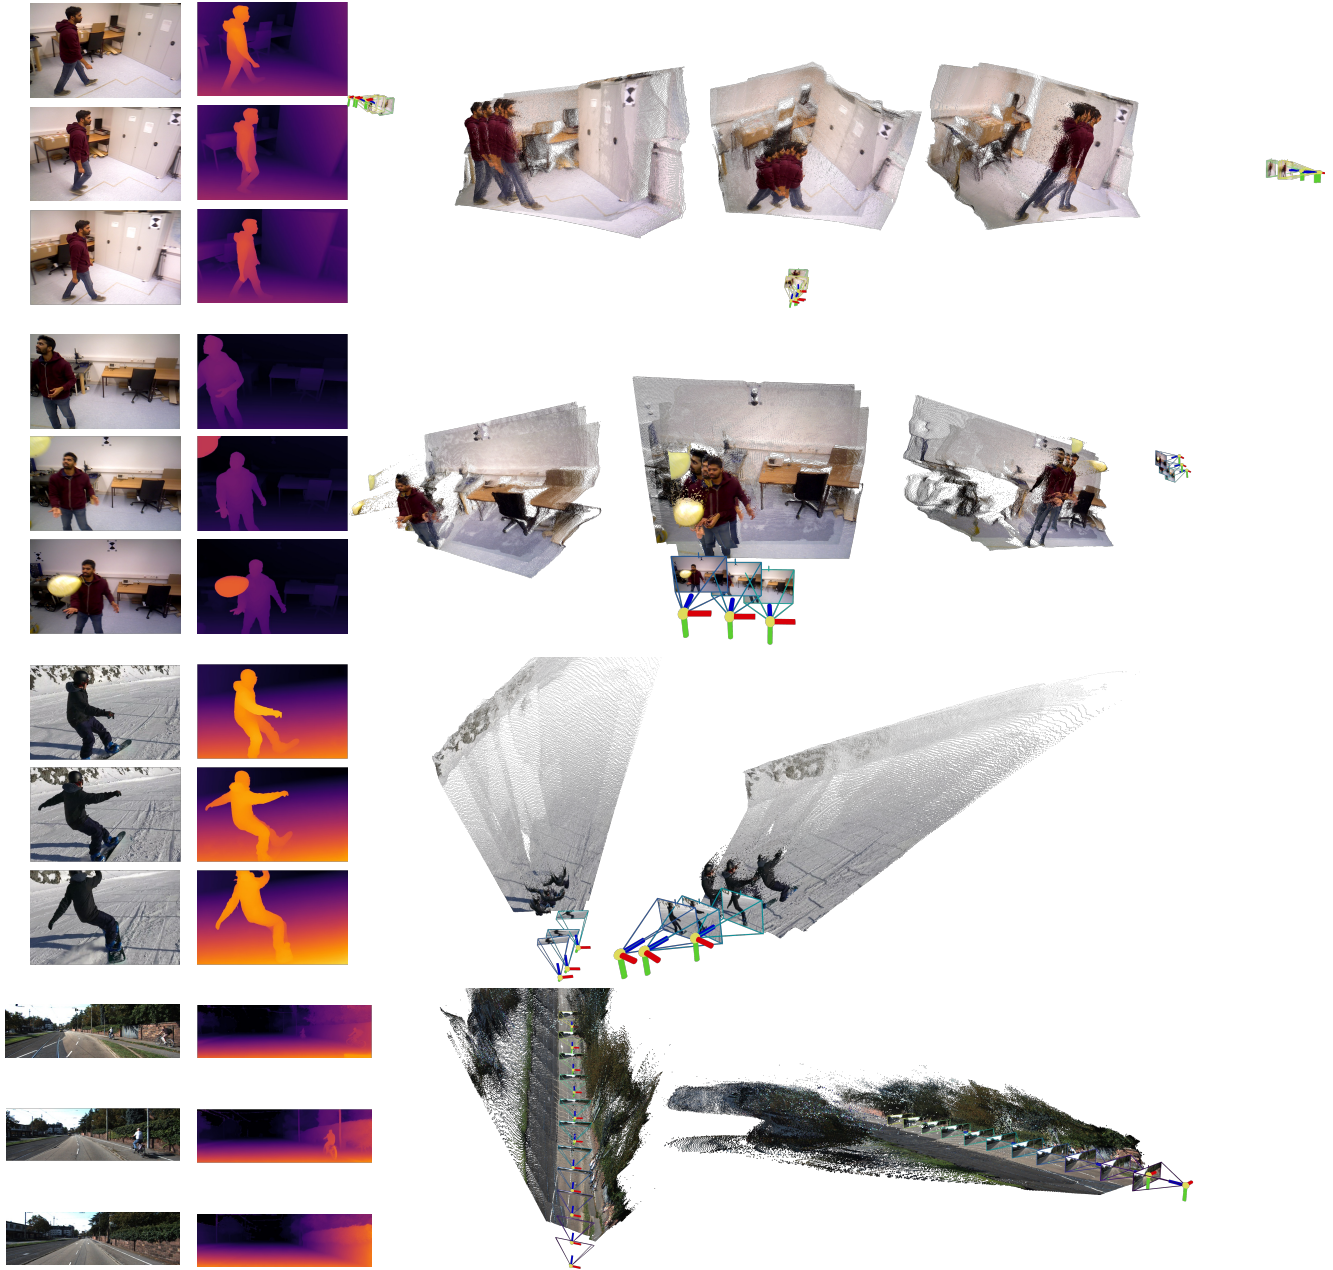


Figure 5. **Additional qualitative results.** Our method generalizes well to various scenes with different 4D objects and performs robustly against different camera and object motions.

Method	Video Depth		Camera Pose		
	Abs Rel ↓	$\delta < 1.25$ ↑	ATE ↓	RPE trans ↓	RPE rot ↓
w/o fine-tuned	0.212	72.1	0.192	0.061	0.577
w fine-tuned	0.205	73.5	0.185	0.063	0.547

Table 7. **Ablation study for the fine-tuned point map VAE** on the Sintel [5] dataset. The fine-tuned point map VAE performs better than the original one.

7.2. Ablation Study for Fine-Tuned Point Map VAE

As stated in the main paper, we added an additional branch to predict the uncertainty for our point map VAE and fine-tuned it based on Eq. 3. We perform an ablation study on our fine-tuning strategy. As shown in Tab. 7, our fine-tuned point map VAE achieves consistently better performance on both video depth estimation and camera pose estimation tasks compared with the original pre-trained image VAE,




Figure 6. Visualization of different geometric modality maps.

demonstrating the necessity and effectiveness of our fine-tuning strategy.

7.3. Analysis of Multi-Modal Representation

Point maps (PMs) and disparity maps (DMs) are complementary. DMs better represent near objects, while PMs are more depth-agnostic (*e.g.*, human vs house in Fig. 6 (b,c)). As in prior work, DMs are affine invariant (which here makes them range-compatible with the pretrained RGB VAE); their scale and shift, needed to recover undistorted geometry, are inferred by matching them to the predicted PMs. Ray maps (RMs) help infer the camera pose when PMs fail to represent points at infinity (such as the sky in Fig. 6 (e)). We observed that PMs tend to be noisier than DMs, so we prioritized modeling the PMs’ uncertainty. Per-pixel uncertainty for ray maps are less meaningful given the high degree of correlation between individual rays. During multi-modal alignment, we align global point clouds with DMs in disparity space and with PMs in linear space. This naturally gives more weight to near points, which tend to be estimated well by DMs, and weighs points based on uncertainty with PMs, thus taking advantage of both modalities.

8. Visualization

Figure 5 shows additional visualizations for indoor, outdoor, and driving scenes. Although our model is only trained on synthetic datasets, it generalizes to real-world data with diverse objects and motions.

9. Limitations

Although our method performs well and generalizes to a wide range of in-the-wild videos, it can struggle in cases involving significant changes in focal length or extreme camera motion throughout a sequence. This limitation likely stems from the lack of focal length variation in our training data. Incorporating more sequences with diverse camera movements and zooming effects could help mitigate this issue. Moreover, due to the inherent temporal attention mechanism in our network architecture, our approach currently supports only monocular video input. Extending the method to handle multi-view images or videos is a promising direction for future work.