

## Video Depth without Video Models

Bingxin Ke<sup>1</sup> Dominik Narnhofer<sup>1</sup> Shengyu Huang<sup>1</sup> Lei Ke<sup>2</sup> Torben Peters<sup>1</sup>  
 Katerina Fragkiadaki<sup>2</sup> Anton Obukhov<sup>1</sup> Konrad Schindler<sup>1</sup>  
<sup>1</sup>ETH Zurich <sup>2</sup>Carnegie Mellon University

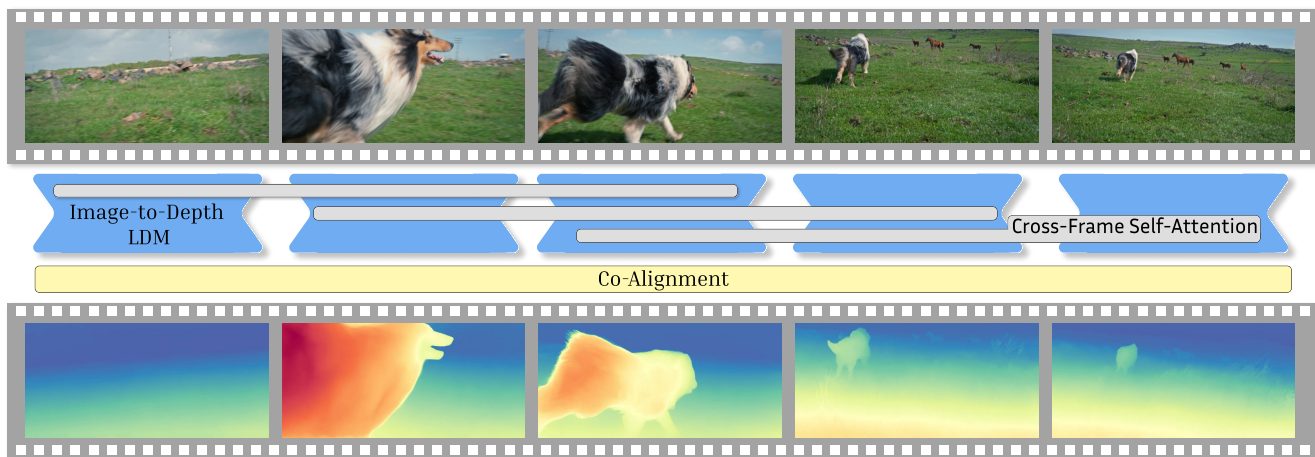


Figure 1. **The RollingDepth model** takes an unconstrained video and reconstructs a corresponding depth video. Unlike methods that rely on video diffusion models, it extends a single-image monodepth estimator such that it can process short snippets. To account for temporal context, snippets with varying frame rates are sampled from the video, processed, and reassembled through a global alignment algorithm to obtain long, temporally coherent depth videos. Depth is colour-coded near far.

### Abstract

Video depth estimation lifts monocular video clips to 3D by inferring dense depth at every frame. Recent advances in single-image depth estimation, brought about by the rise of large foundation models and the use of synthetic training data, have fueled a renewed interest in video depth. However, naively applying a single-image depth estimator to every frame of a video disregards temporal continuity, which not only leads to flickering but may also break when camera motion causes sudden changes in depth range. An obvious and principled solution would be to build on top of video foundation models, but these come with their own limitations; including expensive training and inference, imperfect 3D consistency, and stitching routines for the fixed-length (short) outputs. We take a step back and demonstrate how to turn a single-image latent diffusion model (LDM) into a state-of-the-art video depth estimator. Our model, which we call **RollingDepth**, has two main ingredients: (i) a multi-frame depth estimator that is derived from a single-image LDM and maps very short video snippets (typically frame

triplets) to depth snippets. (ii) a robust, optimization-based registration algorithm that optimally assembles depth snippets sampled at various different frame rates back into a consistent video. RollingDepth is able to efficiently handle long videos with hundreds of frames and delivers more accurate depth videos than both dedicated video depth estimators and high-performing single-frame models. Project page: [rollingdepth.github.io](https://rollingdepth.github.io).

### 1. Introduction

Inferring 3D scene structure from a video stream is a fundamental capability of a vision system. Besides its scientific relevance as an elementary building block of machine perception, it has a broad range of applications, including mobile robotics and autonomous driving, augmented reality, media production, and content creation.

Traditionally, a video would be converted into a 3D world model by recovering the camera trajectory with structure-from-motion (SfM) techniques [16, 53], then ap-

plying multi-view reconstruction based on either stereo triangulation [14, 70] or, more recently, inverse volume rendering [24, 41]. That approach has the attractive property that it delivers a full 3D scene model in a common coordinate frame. The price to pay is that it is only feasible under narrowly defined conditions: the camera motion must be just right, and the scene must have a static background with cooperative texture and lighting conditions. In practice, both SfM and multi-view reconstruction fail more often than not when applied to in-the-wild videos.

This is where video depth comes in. Not all applications require full-scale 3D reconstruction, and it turns out that information about the scene structure can be recovered much more reliably if one aims for a more modest goal: augment every video frame with a dense 2.5D depth map, in such a way that those depth maps are consistent through time. The past years have witnessed tremendous progress in depth estimation from a single image, sidestepping camera pose estimation (and often also calibration of the focal length) [5, 23, 49, 68]. A common thread is that recent methods build on foundation models trained on internet-scale data, such as DINOv2 [43] or StableDiffusion [51], and fine-tune them for depth estimation, often using predominantly synthetic RGB+depth image pairs that can be generated in large quantities and have accurate depth. The underlying, rich visual priors afford these depth estimators excellent zero-shot generalization across scene types, imaging, and lighting conditions.

In general, applying a single-image depth estimator to a video frame-by-frame does not yield satisfactory results, but leads to depth flicker and drift. These artifacts are caused by multiple factors. Most obviously, neither the model training nor the inference procedure have any notion of temporal coherence between adjacent frames. Moreover, monodepth estimation requires scene understanding, which may also suffer from the lack of temporal context (e.g., when a partially visible object only becomes recognizable after zooming out). What is more, in a video the depth range between nearby and distant scene parts may change all of a sudden (e.g., when a foreground object enters the viewfield, or when the camera pans to a window), making consistent monodepth estimation difficult.

Some authors [22, 55] have explored the idea of repurposing generative video models like Stable Video Diffusion [4] for depth prediction. These methods enable information exchange along the time axis and acquire a strong flow and motion prior during training, hence they achieve excellent local consistency through time. On the downside, video LDMs – besides being computationally demanding – are trained for fixed, short sequence lengths and cannot be applied directly to uncurated footage of varying lengths. To be practically useful, the diffusion routine must be wrapped into a partitioning scheme that splits the input video for pro-

cessing and stitches the depth estimates back together, often resulting in low-frequency flickering and gradual drift. We also find that current LDM-based video depth models tend to be less accurate on distant scene parts.

Rather than design more refined video LDMs, which require huge resources for training, we take a step back and re-examine how far one can take video depth estimation with augmented *single-image* LDMs. We design a set of measures that, taken together, extend a per-image monodepth framework like Marigold [23] in a way that enables it to handle video input. Importantly, these measures greatly improve local and global consistency across time while maintaining a constant memory footprint such that one can process long sequences. Specifically, we employ a “rolling” inference with a sliding window of a few frames (typically three, but other numbers are possible). Those snippets are sampled from the video with varying spacing, i.e., they can be immediately adjacent but also dilated along the timeline to cover long-range context. They are then fed into a multi-frame LDM fine-tuned from a single-frame model, with a modified cross-frame self-attention mechanism to enable information exchange. To reassemble the snippets, we propose a robust optimization-based global co-alignment, followed by averaging the aligned frames. Optionally, the resulting video can be degraded with moderate random noise and denoised again with the same per-snippet LDM to further refine spatial details.

To summarize, our approach estimates accurate and temporally consistent video depth without resorting to cumbersome video diffusion models. To that end, we contribute:

1. an LDM for monocular depth estimation in video snippets of a few frames, adapted from the Marigold [23] single-frame model but able to capture temporal patterns across frames via self-attention;
2. a rolling inference scheme that operates on snippets with multiple different (temporal) resolutions and enables efficient propagation of contextual information through video sequences of arbitrary length (up to minutes);
3. a global alignment procedure, based on robust optimization, to recompose the snippets into a depth video whose depth values remain consistent over long time periods;
4. an optional refinement of the final output with another round of multi-frame diffusion, where the same LDM is applied starting from a moderately degraded video.

## 2. Related Work

### 2.1. Monocular Depth Estimation

Monocular depth estimation is a dense regression task. The pioneering work by Eigen et al. [11] showed that metric depth values can be recovered from single sensors. Successive advancements include including various parameterizations (ordinals, bins, planar guidance maps, piecewise pla-

narity, CRFs, *etc.*) [2, 12, 29, 31, 36, 42, 45, 75, 81], switching CNN backbones to vision transformers [1, 3, 32, 65], considering camera intrinsics [18, 21, 46, 47, 74], and patch-wise processing [5, 34, 35]. To handle “in-the-wild” settings, extensive internet photo collections are used for training [30, 73]. MiDaS [49] improves the generality by training on a mixture of multiple datasets. Depth Anything [67, 68] takes data scaling to the next level by relying on DINOv2 [43], a foundational model trained on 142M images in a self-supervised manner, and subsequently training with 62M pseudo-labels, 1M real depth annotations, and 0.5M synthetic ones. Recent trends leverage generative models, particularly diffusion models [20, 56], for depth estimation [10, 52, 80, 80]. Marigold [23] proposed to fine-tune Stable Diffusion [51], a generative text-to-image latent diffusion model (LDM) trained with LAION-5B [54], towards affine-invariant depth using 74k samples. This approach has been improved in many aspects including fewer steps [15, 17, 19, 64], finer details [77], and more modalities [13, 19].

## 2.2. Video Depth Estimation

Video depth estimation calls for dedicated mechanisms to ensure smoothness of adjacent frames, and correct handling of varying depth range. Existing approaches can be grouped into three main categories: test-time optimization, feed-forward prediction, and diffusion-based. *Test-time optimization* methods [7, 27, 40, 78] often rely on camera poses or optical flow and perform optimization for each new video during inference. While these methods can produce depth estimates that are temporally consistent, their dependence on camera poses and long processing time hamper their application to open-world video scenarios. *Feed-forward* prediction methods estimate depth sequences directly from input videos [33, 58, 60, 71, 72, 76]. For example, DeepV2D [58] integrates camera motion estimation with depth prediction, MAMO [71] adopts memory attention mechanisms, and NVDS [61, 62] introduces a stabilization network as a post-processing module. However, the generalization of these methods to in-the-wild videos is often constrained by the limited diversity of training data and model capacity.

Very recently, concurrent with our work, several authors have investigated the use of video diffusion models, in particular SVD [4], for video depth. ChronoDepth [55] DepthCrafter [22] and DepthAnyVideo [66] all modify video diffusion for conditional generative depth prediction. From the underlying video diffusion model they inherit high training and inference costs, and a restriction to short video clips of at most  $\approx 100$  frames. In contrast, in RollingDepth we explore how to turn an image diffusion model into a temporally consistent depth estimator, which can handle long videos of 1000 frames or more.

## 2.3. Image Diffusion Models for Video Tasks

Image diffusion models have been employed in various video inverse problems, such as video generation, inpainting, and super-resolution [9, 28, 82]. A large amount of work [48, 69, 79] focusses on video editing, either by fine-tuning text-to-image diffusion models on video data [37, 63] or through training-free approaches using cross-frame attention and latent fusion [6, 25]. However, these works [6, 37, 63] predominantly address video-to-video translation tasks, where both the input and output reside in RGB space. In contrast, our approach leverages image diffusion priors to generate consistent depth videos, with the additional challenge to accommodate large variations of the depth range, as the near and far planes change – often suddenly – due to camera and object motion. Implementation tricks when using single-image models, like fixing the initial noise or blending consecutive latent representations, can somewhat mitigate the lack of knowledge w.r.t. temporal coherence, but do not solve it [23].

## 3. Method

Let  $\mathbf{x} \in \mathbb{R}^{N_F \times 3 \times H \times W}$  be an RGB video of length  $N_F$ , the goal of a monocular video depth estimator is to predict a depth video  $\mathbf{d} \in \mathbb{R}^{N_F \times H \times W}$ . All frames in that depth video should share a common depth scale and shift, i.e., depth values should not drift unless the associated pixel moves relative to the camera. In the following, we present our RollingDepth framework for predicting  $\mathbf{d}$  from  $\mathbf{x}$ . The proposed approach is based on a per-snippet LDM, test-time depth co-alignment, and an optional refinement of the resulting video, as illustrated in Fig. 2.

### 3.1. Marigold Monocular Depth Recap

Several recent methods [10, 23, 52], including our base model Marigold [23], cast monocular depth estimation as conditional image generation, where a pre-trained LDM is retargeted to generate the depth map given the input image. To that end, the model progressively adds noise to depth samples  $\mathbf{d}^i$  and learns to reverse that degradation, to approximate the conditional distribution  $p(\mathbf{d}^i | \mathbf{x}^i)$ .

In detail, the model is trained to predict the added noise  $\epsilon$  at each step by minimizing the objective

$$\mathcal{L}(\theta) = \mathbb{E}_{(\mathbf{d}_0^i, \mathbf{x}^i) \sim P_{\mathbf{d}^i, \mathbf{x}^i}, t \sim \mathcal{U}, \epsilon \sim \mathcal{N}} \left[ \left\| \epsilon - \epsilon_\theta(\mathbf{d}_t^i, \mathbf{x}^i, t) \right\|^2 \right].$$

At inference time the model starts from the input  $\mathbf{x}^i$  and pure Gaussian noise  $\mathbf{d}_T^i \sim \mathcal{N}(0, \mathbf{I})$ , and gradually maps the latter to a depth map  $\mathbf{d}_0^i$  by iteratively applying the learned denoising step. For computational efficiency, the denoising process operates a low-dimensional latent space  $\mathcal{Z}$ , with an auto-encoder to map images to latent embeddings, and depth maps back to image space [51].

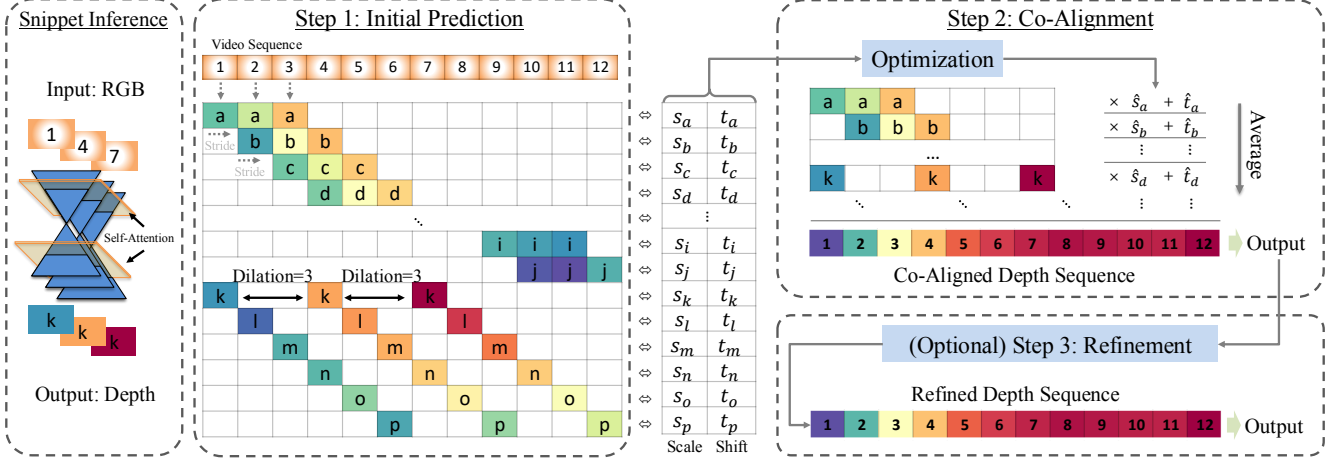


Figure 2. **Overview of the RollingDepth Inference Pipeline.** Given a video sequence  $\mathbf{x}$  (with  $\mathbf{x}^i$  is  $i^{\text{th}}$  frame), we construct  $N_T$  overlapping snippets using a *dilated rolling kernel* with varying dilation rates, and perform 1-step inference to obtain initial depth snippets (  $\begin{bmatrix} k & k & k \end{bmatrix}$  ). Next, depth co-alignment optimizes  $N_T$  pairs of scale and shift values to achieve globally consistent depth throughout the full video. An optional refinement step further enhances details by applying additional, snippet-based denoising steps.

### 3.2. Extension to Snippets

Inspired by multi-view diffusion models [26, 38], we extend Marigold [23] to handle multiple frames by modifying its self-attention layers. In each self-attention block, we flatten tokens from all frames in a snippet into a single sequence, such that the attention mechanism operates across frames and captures spatial and temporal interactions. Unlike video diffusion models with factorized spatial-temporal attention, this approach can handle frames with varying temporal spacing, which makes it possible to sample snippets at lower frame rates and capture long-range dependencies, an advantage when processing long videos.

The original Marigold model predicts (affine-invariant) depth between image-specific near and far planes. This parametrization poses problems for video depth estimation, where the depth range can vary over time. We therefore re-train Marigold to predict inverse depth (like several other monodepth estimators [49, 68]), which is less sensitive to such variations, particularly in the far field.

### 3.3. From Snippets to Video

Our multi-frame depth estimator operates on short snippets of  $n$  frames, where  $n \ll N_F$ . As these snippets are processed independently, each has its own scale and shift – which are arbitrary in the case of affine-invariant methods [17, 23, 49] including Marigold, but will in practice not be perfectly aligned even when using a metric depth estimator [35, 46, 74]. To resolve that ambiguity, we construct overlapping snippets with different temporal dilation rates. The frames shared between different snippets are subsequently used to align all depth predictions to one common scale and shift.

**Dilated Rolling Kernel.** We construct multi-scale snippets using the *dilated rolling kernel*. For instance, for 3-frame snippets with dilation rate (frame spacing)  $g$  and stride  $h$ , the kernel picks frames  $(\mathbf{x}^{i-g}, \mathbf{x}^i, \mathbf{x}^{i+g})$  from the input video, where  $i \in \{g+1, g+1+h, g+1+2h, \dots\}$ . By varying the dilation rate, we sample snippets with different frame rates, in order to capture temporal dependencies at different time scales. For each snippet of  $n$  frames, we then predict depth using the multi-frame LDM, to obtain a corresponding  $n$ -frame depth snippet.

**Depth Co-alignment.** At this stage we have generated  $N_T$  depth snippets. Each of them has its own scale and shift parameters  $\{(s_k, t_k), k \in 1 \dots T\}$ , which are shared across its constituent frames. Our goal is to jointly compute  $N_T$  scale and shift values such that they optimally align all snippets into a consistent video. At a given frame  $\mathbf{x}^i$ , there are  $N^i$  different individual depth maps  $\{\mathbf{d}_j^i, j = 1 \dots N^i\}$  originating from different snippets, where  $N^i$  can vary from frame to frame. Let  $k(i, j)$  be an indexing function that retrieves the snippet index  $k$  for the  $j$ -th depthmap at frame  $i$ . To estimate the best alignment, we minimize the  $L_1$  loss over all individual depth predictions,

$$\min_{s_k > 0, t_k} \left( \sum_{i=1}^{N_F} \sum_{j=1}^{N^i} |s_{k(i,j)} \mathbf{d}_j^i + t_{k(i,j)} - \bar{\mathbf{d}}^i| \right), \quad (1)$$

with the mean depth

$$\bar{\mathbf{d}}^i = \frac{1}{N^i} \sum_{j=1}^{N^i} (s_{k(i,j)} \mathbf{d}_j^i + t_{k(i,j)}) . \quad (2)$$

The solution to eq. (1) is found with gradient descent, stabilized by putting more emphasis on snippets with high di-



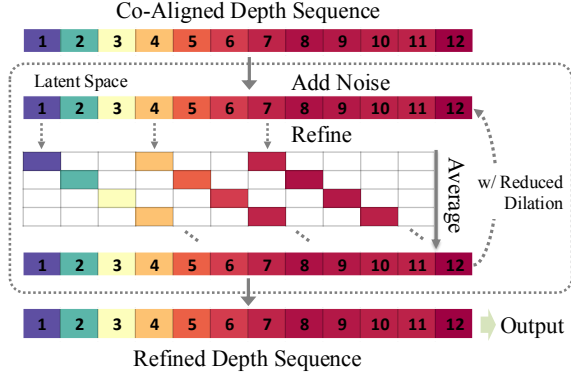


Figure 3. **Depth Refinement** encodes the co-aligned depth video into latent space, contaminates it with a moderate amount of noise, then denoises it with a series of reverse diffusion steps with decreasing snippet dilation rate. After each step, overlapping latents are averaged to propagate information between snippets.

lution rates, and additional regularization. Once the depth snippets have been aligned in a common frame with the estimated scale and shift values, the depth maps at every frame  $\mathbf{x}^i$  are obtained by taking the pixel-wise mean  $\bar{\mathbf{d}}^i$ , resulting in a single, consistent depth video with one depth map per frame. See the supplementary material for further details.

**Depth Refinement.** To enhance visual quality and capture finer details, we optionally apply a diffusion-based refinement step to the merged depth video  $\mathbf{d}$ , as illustrated in Fig. 3. The video is again encoded into latent space frame by frame, and contaminated with a moderate amount of noise, corresponding to step  $T/2$  of the diffusion schedule, halfway between the clean latent and pure noise. The degraded video is again split into snippets with the dilated rolling kernel and each snippet is denoised individually with the same LDM as above. To integrate information across overlapping snippets, the latent embeddings of every frame are averaged after every denoising step. We find that this partial (reverse) diffusion works best when applied in a coarse-to-fine manner in time, starting with a large snippet dilation rate and gradually decreasing it along the denoising process. The refinement process enhances high-frequency detail without altering the global scene depth layout, at the cost of increased inference time due to the additional round of denoising diffusion.

### 3.4. Multi-Frame Training

We exploit the flexible design of the multi-frame self-attention mechanism to fine-tune the model with varying snippet lengths. Training snippets are randomly picked to have one, two, or three frames, making sure that the motion between frames is small enough to have overlapping view frustra. To fully utilize the value range of the diffusion model for best performance, inverse depth values are nor-

malized on a per-snippet basis, using the 2<sup>nd</sup> and 98<sup>th</sup> percentiles for robustness. We found it important to jointly normalize the values within each snippet rather than normalizing each frame individually. In this way, the same frame is normalized differently depending on the context it appears in, and normalized depths remain comparable within a snippet, enabling the model to understand and correctly handle rapid changes in the depth range, which routinely appear in longer video sequences.

## 4. Experiments

### 4.1. Implementation Details

**Training Datasets.** To finetune the snippet LDM, we use TartanAir [59], a synthetic video dataset with various (indoor and outdoor) scenes, styles, and camera motions. We visually inspect the scenes and select 18 scenes consisting of 369 sequences. Training snippets are randomly sampled from a sequence, with a minimum overlap ratio of 30%. To increase the diversity of scenes and avoid a significant sim-to-real gap we additionally use Hypersim [50], a photorealistic single-image dataset containing 365 diverse scenes, treating images as 1-frame snippets.

**Training Settings.** Training images are resized to  $480 \times 640$  for efficiency, with random horizontal flipping as data augmentation. To align with the refinement setting, we employ depth range augmentation, where we randomly squeeze the normalized depth snippets to a smaller range and then slightly rescale and shift the depth range in each frame. As optimizer, we use AdamW [39] with a learning rate of  $3 \times 10^{-5}$  and exponential decay. Training is run on four Nvidia A100 GPUs with a batch size of 32 and takes approximately 18k iterations or two days to converge.

**Inference Settings.** During inference, we fix the snippet length to  $n = 3$ , with three different dilation rates  $g \in \{1, 10, 25\}$  to capture short- to mid-range temporal relations. For each snippet we perform 1-step inference. Long-range temporal relations are covered by the depth co-alignment, which is initialized with  $s_k = 1$  and  $t_k = 0$  and optimized with 2000 steps of gradient descent, using the Adam optimizer. For the optional refinement, we start at timestep  $T/2$  of the diffusion trajectory and perform 10 denoising steps, gradually reducing the dilation rate from 6 to 1. Input images are resized to a maximum side length of 768 pixels. For evaluation, the final result is up-sampled to match the original resolution in the dataset.

### 4.2. Evaluation

**Evaluation Datasets.** We evaluate RollingDepth on four datasets that include both static and dynamic scenes with varying camera and scene motions: **PointOdyssey** [83] is a synthetic dataset with individually animated characters that move independently, designed for long-term tracking.

Table 1. **Quantitative comparison** of RollingDepth with baseline methods on zero-shot benchmarks. **Bold** numbers are the best, underscored second best, numbers in the bracket after each dataset denote video sequence length. RollingDepth demonstrates superior performance across both short and long video sequences, despite being an image-based model.

		PointOdyssey (250)		ScanNet (90)		Bonn (110)		DyDToF (200)		DyDToF (100)	
		Abs Rel↓	$\delta 1 \uparrow$	Abs Rel↓	$\delta 1 \uparrow$	Abs Rel↓	$\delta 1 \uparrow$	Abs Rel↓	$\delta 1 \uparrow$	Abs Rel↓	$\delta 1 \uparrow$
Single frame	Marigold* [23]	14.9	80.4	14.9	78.3	10.5	86.7	25.3	55.5	16.4	73.5
	DepthAnything [67]	16.3	76.0	12.9	84.0	9.9	89.4	25.4	54.3	16.4	75.6
	DepthAnythingv2 [68]	14.4	81.4	13.3	82.6	10.5	87.4	24.8	55.9	16.0	76.6
Video	NVDS (DPT-Large) [61]	26.6	68.2	18.5	67.7	10.5	88.1	24.7	56.0	18.8	69.3
	ChronoDepth [55]	51.7	71.2	16.8	73.8	10.9	86.9	26.9	53.2	19.9	66.5
	DepthCrafter [22]	36.3	75.0	12.7	84.3	<b>6.6</b>	<b>96.7</b>	22.1	60.7	16.2	74.7
	RollingDepth (ours, fast) <sup>†</sup>	<b>9.6</b>	<u>90.4</u>	<u>10.1</u>	<u>89.7</u>	<u>7.9</u>	93.6	<u>17.7</u>	<u>69.6</u>	<u>12.7</u>	<u>81.6</u>
	RollingDepth (ours)	<b>9.6</b>	<b>90.5</b>	<b>9.3</b>	<b>91.6</b>	<u>7.9</u>	<u>93.9</u>	<b>17.3</b>	<b>71.7</b>	<b>12.3</b>	<b>83.0</b>

\* Inverse depth version, retrained with the original training code. <sup>†</sup> Run at half-precision (fp16), with dilation rates {1, 25}, without refinement.

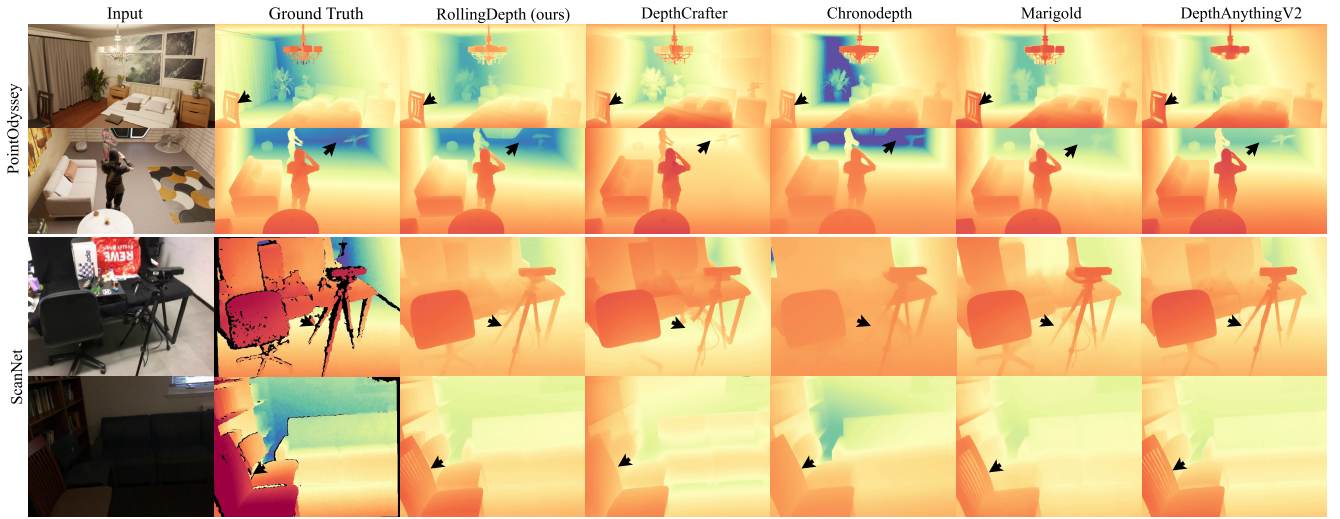


Figure 4. **Qualitative comparison** between different methods. RollingDepth excels at preserving fine-grained details (*cf.* the chandelier in the first sample and the tripod in the third sample) and recovering accurate scene layout (*cf.* the far plane in the second sample).

We filter out overly simplistic toy scenes and retain 35 sequences. For each sequence, we follow the videodepth literature and exclude frames with camera zoom, then select the first 250 frames in each sequence. **ScanNet v2** [8] is an indoor dataset of static scenes recorded with the Kinect RGB-D sensor. We use its test set of 100 sequences, taking the first 270 frames of each sequence and downsample the frame rate by a factor 3, **Bonn RGBD** [44] is an RGB-D dataset of moving people in indoor spaces. Following [22], we use frames 30-140 from five different dynamic scenes. **DyDToF** [57] is a photorealistic synthetic dataset featuring moving objects including people and animals. It has several videos per scene, we always take the first video and create two subsets of different lengths from it, by clipping frames 50-250, respectively frames 50-150.

**Evaluation Protocol.** We extend the affine-invariant depth evaluation protocol [49] to videos, i.e., depth predictions  $\hat{\mathbf{d}}$  are aligned to the ground truth with a scale and shift found

with least squares fitting, where we fit one pair of transformation parameters per video, i.e., all frames in a video share a common scale and shift. We quantify the depth estimation accuracy with two standard metrics [22, 23, 49, 61]: the *absolute mean relative error* (*AbsRel*), defined as  $\frac{1}{M} \sum_{j=1}^M |\hat{\mathbf{d}}_j - \mathbf{d}_j| / \mathbf{d}_j$ , where  $M$  is the total number of pixels; and the  $\delta 1$ -accuracy, which measures the fraction of pixels for which  $\max(\hat{\mathbf{d}}_j / \mathbf{d}_j, \mathbf{d}_j / \hat{\mathbf{d}}_j) < 1.25$ . Metrics are always given as percentages. We provide additional temporal smoothness evaluation in the supplementary material.

### 4.3. Comparison with Other Methods

We compare RollingDepth against six state-of-the-art methods for zero-shot monocular depth estimation: Marigold [23], DepthAnything [67] and DepthAnythingv2 [68], which are single-frame methods; as well as NVDS [62], ChronoDepth [55], and DepthCrafter [22], which are video-based approaches.

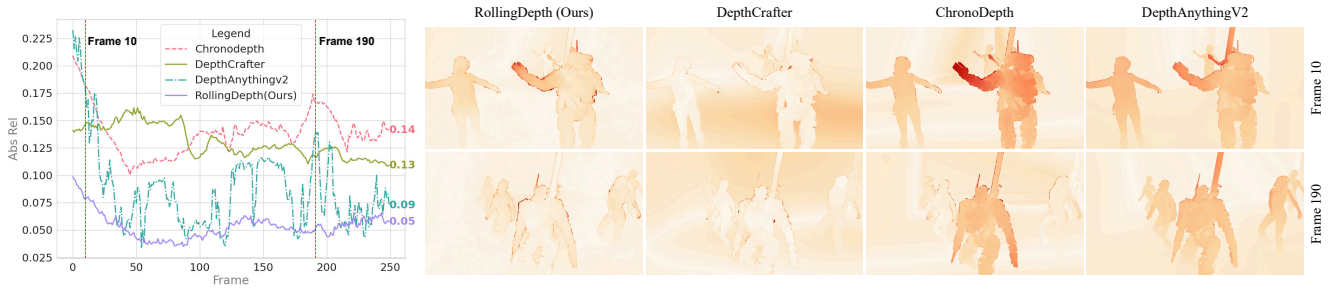


Figure 5. **AbsRel error over time:** The line plot (left) shows the depth error at every individual frame, end-of-line numbers are the average error across the video. The images (right) display error maps (low  $\rightarrow$  high) for two specific frames. RollingDepth achieves the lowest error overall, competing methods recover scene layout less faithfully and tend to be biased towards the foreground or the background.

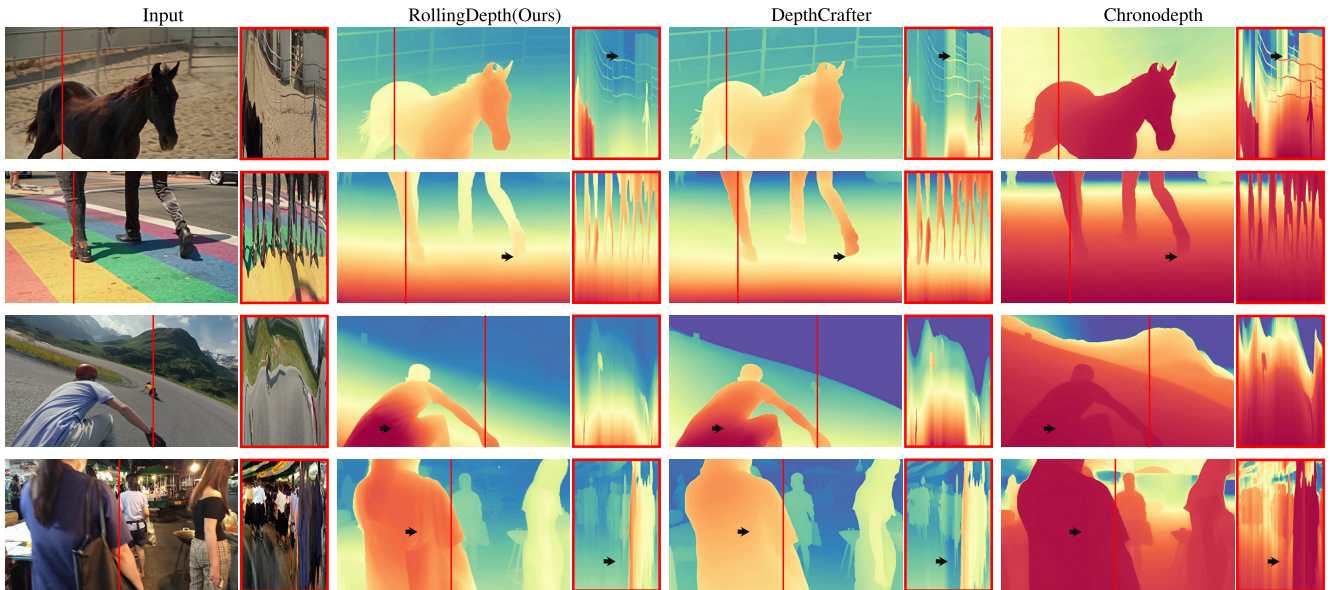


Figure 6. **Qualitative comparison** of depth predictions (near  $\rightarrow$  far) from in-the-wild videos. To graphically show temporal consistency, we display temporal profiles (red box) for a fixed column (marked with a red line). RollingDepth picks up subtle details like accessories and wrinkled cloth, and mitigates spurious depth discontinuities (cf. background in temporal profile of the first sample) in time.

**Quantitative Comparison.** As shown in Tab. 1, RollingDepth outperforms both single-frame and video-based approaches across multiple datasets and different sequence lengths, often by considerable margins. We attribute this to its ability to combine the accuracy of image-based models with the temporal coherence afforded by our snippet-based inference and global depth co-alignment. On PointOdyssey, which includes many challenging scenes with highly variable depth ranges, RollingDepth achieves by far the best result. Methods based on video models struggle on this dataset, and are in fact even unable to match the performance of single-frame methods. We observe that the performance of video models drops especially in scenes with sudden, large changes in the depth range (e.g. a hand gesture in front of the camera). We hypothesize that the underlying video prior is too rigid and prevents a correct adaptation to the rapid change, see the supplementary material

for details. Also on DyDTof, RollingDepth greatly reduces the error compared to other methods, again underscoring its ability to handle dynamic scenes and variations of the near and far planes. Still, the good performance is not limited to dynamic scenes with strong depth variations. RollingDepth also performs well on indoor data, reaching the lowest error on the static ScanNet scenes and the second-lowest error on the Bonn data. Here DepthCrafter shines – we observe that it generally tends to do well in scenes dominated by foreground objects, particularly humans.

**Qualitative Comparison.** To make our findings more tangible, we provide qualitative comparisons both on evaluation data and on in-the-wild examples. Figure. 4 confirms that RollingDepth consistently produces high-quality depth maps that preserve fine detail, both near the camera in the distance. DepthCrafter and ChronoDepth produce locally smooth videos with little frame-to-frame flicker, but have a



tendency to distort the overall scene layout in a way that certain objects are segmented well but placed at incorrect (relative) depths. Single-image estimators are seemingly more accurate in that respect, but suffer from flickering and a lack of temporal coherence. We further illustrate these trends in Fig. 5, where we plot per-frame errors, as well as per-pixel errors for selected frames.

To demonstrate generalization to real-world video clips, Fig. 6 shows depth predictions for videos collected from the internet. Also in these cases, RollingDepth accurately recovers fine details and maintains long-term coherence. To better illustrate the evolution of the depth estimate over time, we extract temporal profiles for fixed image columns. They exhibit no significant high-frequency variations along the time axis that would indicate frame-to-frame flicker. We also do not observe drift or unwarranted jumps in the depth values that would indicate systematic biases. DepthCrafter for the most part also recovers plausible depth, but misses depth variations within the main segments and sometimes exhibits instabilities along the time axis. Chronodepth recovers depth boundaries rather well, but delivers billboard-like, layered depth maps.

#### 4.4. Ablation Studies

We validate our main hyper-parameters and design choices on a subset of 10 sequences from the PointOdyssey test set and 20 sequences from the ScanNet test set.

**Dilation of Initial Predictions.** We start by ablating the arguably most crucial hyper-parameter of RollingDepth, the dilation rate for snippet sampling, see Tab. 2. The base setting uses only dilation rate  $\{1\}$  for minimal information exchange and smoothness between adjacent frames. Having a high dilation rate  $\{1, 25\}$  gives the model access to longer-term motion patterns on the order of 1 second and greatly stabilizes the co-alignment step, which in turn reduces the AbsRel error by  $>6$  percept points on PointOdyssey and by  $>2$  percent points on the (static) ScanNet. This is what we use in our *fast* setting (c.f. Tab. 1), which takes 81s for a  $768 \times 432$  video of 250 frames (ChronoDepth: 121s, DepthCrafter: 284s). An additional, intermediate dilation rate  $\{1, 10, 25\}$  further intensifies the information exchange across time. This further boosts the quality of the estimated depth maps, but as expected yields diminishing returns.

**Effectiveness of Co-Alignment and Refinement.** We further isolate the effect of the RollingDepth’s components, see Tab. 3. The snippet diffusion step is mandatory to obtain any depth estimates at all and cannot be left out. For the experiment we switch on and off the two remaining steps, co-alignment and refinement, and test all combinations. Simply merging overlapping latents without prior alignment proves to be insufficient, i.e., their individually estimated depth ranges are too inconsistent to average them into a coherent sequence. The refinement step cannot fix

Table 2. **Ablation of dilation rates** for snippet prediction. We report values before the optional refinement step. The minimal base setting uses only dilation rate 1. Adding a high dilation rate 25 brings a marked performance gain. Yet another dilation rate 10 gives a further, smaller boost.

Dilation rates	PointOdyssey		ScanNet	
	Abs Rel $\downarrow$	$\delta 1 \uparrow$	Abs Rel $\downarrow$	$\delta 1 \uparrow$
$\{1\}$	16.7	75.5	12.8	83.2
$\{1, 25\}$	<b>10.2</b>	89.5	10.6	88.8
$\{1, 10, 25\}$	<b>10.2</b>	<b>89.6</b>	<b>9.9</b>	<b>90.1</b>

Table 3. **Ablation of components.** Depth co-alignment is a crucial functionality for the snippet-based strategy of RollingDepth, whereas the additional refinement has only a small effect on the performance metrics, despite visibly enhanced image detail.

Co-Alignment	Refinement	PointOdyssey		ScanNet	
		Abs Rel $\downarrow$	$\delta 1 \uparrow$	Abs Rel $\downarrow$	$\delta 1 \uparrow$
$\times$	$\times$	13.0	84.4	12.4	84.3
$\times$	$\checkmark$	13.0	84.6	12.3	84.8
$\checkmark$	$\times$	<b>10.2</b>	89.6	9.9	90.1
$\checkmark$	$\checkmark$	<b>10.2</b>	<b>89.8</b>	<b>9.8</b>	<b>90.2</b>

that problem. Conversely, the co-alignment does the heavy lifting to fuse depth snippets with different scales and shifts into a coherent video and contributes the lion’s share of the improvement. Subsequent refinement of the aligned video only results in a marginal increase of the performance metrics, but visibly improves the result by recovering sharp details that have been missed or blurred in the preceding steps.

## 5. Conclusion

We have introduced RollingDepth, a novel method for monocular video depth estimation that is derived from a single-image (latent) diffusion model. The core components of our method (i) are a monodepth estimator for short snippets, sampled at various frame rates to capture temporal context at different time scales; (ii) an optimization-based co-alignment procedure that optimally registers all snippets of a video into a common depth range; and (iii) an optional refinement step, again based on the same denoising diffusion scheme for snippets, that enhances fine details in the depth video. RollingDepth strikes a favorable balance between accurate per-frame depth prediction and temporal coherence, and can process long video with hundreds of frames. It empirically delivers best-in-class performance across multiple datasets, also outperforming alternatives derived from full-blown video diffusion models. That being said, the RollingDepth framework is flexible and offers the possibility to replace individual components. For instance, an interesting avenue for future work would be to swap out the snippet-based refinement and replace it with a generative video model or a flow-based method for even better motion reconstruction.



**Acknowledgments.** We thank Yue Pan, Shuchang Liu, Nando Metzger, and Nikolai Kalischek for fruitful discussions. We are grateful to *redmond.ai* for providing GPU resources.

## References

- [1] Shubhra Aich, Jean Marie Uwabeza Vianney, Md Amirul Islam, Mannat Kaur, and Bingbing Liu. Bidirectional attention network for monocular depth estimation. In *ICRA*, 2021. 3
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. AdaBins: Depth estimation using adaptive bins. In *CVPR*, 2021. 3
- [3] Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. ZoeDepth: Zero-shot transfer by combining relative and metric depth. *arXiv:2302.12288*, 2023. 3
- [4] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, 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:2311.15127*, 2023. 2, 3
- [5] Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R Richter, and Vladlen Koltun. Depth Pro: Sharp monocular metric depth in less than a second. *arXiv:2410.02073*, 2024. 2, 3
- [6] Duygu Ceylan, Chun-Hao P Huang, and Niloy J Mitra. Pix2Video: Video editing using image diffusion. In *ICCV*, 2023. 3
- [7] Yuhua Chen, Cordelia Schmid, and Cristian Sminchisescu. Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In *ICCV*, 2019. 3
- [8] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Niessner. ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. 6
- [9] Giannis Daras, Weili Nie, Karsten Kreis, Alex Dimakis, Morteza Mardani, Nikola Borislavov Kovachki, and Arash Vahdat. Warped diffusion: Solving video inverse problems with image diffusion models. *NeurIPS*, 2024. 3
- [10] Yiqun Duan, Xianda Guo, and Zheng Zhu. DiffusionDepth: Diffusion denoising approach for monocular depth estimation. *arXiv:2303.05021*, 2023. 3
- [11] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *NeurIPS*, 2014. 2
- [12] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, 2018. 3
- [13] 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 *ECCV*, 2024. 3
- [14] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015. 2
- [15] Gonzalo Martin Garcia, Karim Abou Zeid, Christian Schmidt, Daan de Geus, Alexander Hermans, and Bastian Leibe. Fine-tuning image-conditional diffusion models is easier than you think. *arXiv:2409.11355*, 2024. 3
- [16] Carsten Griwodz, Simone Gasparini, Lilian Calvet, Pierre Gurdjos, Fabien Castan, Benoit Maujean, Gregoire De Lillo, and Yann Lanthony. AliceVision Meshroom: An open-source 3d reconstruction pipeline. In *ACM Multimedia*, 2021. 1
- [17] Ming Gui, Johannes S Fischer, Ulrich Prestel, Pingchuan Ma, Dmytro Kotovenko, Olga Grebenkova, Stefan Andreas Baumann, Vincent Tao Hu, and Björn Ommer. DepthFM: Fast monocular depth estimation with flow matching. *arXiv:2403.13788*, 2024. 3, 4
- [18] Vitor Guizilini, Igor Vasiljevic, Dian Chen, Rareş Ambruş, and Adrien Gaidon. Towards zero-shot scale-aware monocular depth estimation. In *ICCV*, 2023. 3
- [19] Jing He, Haodong Li, Wei Yin, Yixun Liang, Leheng Li, Kaiqiang Zhou, Hongbo Liu, Bingbing Liu, and Ying-Cong Chen. Lotus: Diffusion-based visual foundation model for high-quality dense prediction. *arXiv:2409.18124*, 2024. 3
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 3
- [21] Mu Hu, Wei Yin, Chi Zhang, Zhipeng Cai, Xiaoxiao Long, Hao Chen, Kaixuan Wang, Gang Yu, Chunhua Shen, and Shaojie Shen. Metric3D v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *arXiv:2404.15506*, 2024. 3
- [22] Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiaodong Cun, Yong Zhang, Long Quan, and Ying Shan. DepthCrafter: Generating consistent long depth sequences for open-world videos. *arXiv:2409.02095*, 2024. 2, 3, 6
- [23] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *CVPR*, 2024. 2, 3, 4, 6
- [24] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d Gaussian splatting for real-time radiance field rendering. *ACM TOG*, 42(4):139–1, 2023. 2
- [25] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2Video-Zero: Text-to-image diffusion models are zero-shot video generators. In *ICCV*, 2023. 3
- [26] Xin Kong, Shikun Liu, Xiaoyang Lyu, Marwan Taher, Xiaojuan Qi, and Andrew J Davison. EscherNet: A generative model for scalable view synthesis. In *CVPR*, 2024. 4
- [27] Johannes Kopf, Xuejian Rong, and Jia-Bin Huang. Robust consistent video depth estimation. In *CVPR*, 2021. 3
- [28] Taesung Kwon and Jong Chul Ye. Solving video inverse problems using image diffusion models. *arXiv:2409.02574*, 2024. 3
- [29] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv:1907.10326*, 2019. 3
- [30] Zhengqi Li and Noah Snavely. MegaDepth: Learning single-view depth prediction from internet photos. In *CVPR*, 2018. 3

- [31] Zhenyu Li, Xuyang Wang, Xianming Liu, and Junjun Jiang. BinsFormer: Revisiting adaptive bins for monocular depth estimation. *arXiv:2204.00987*, 2022. 3
- [32] Zhenyu Li, Zehui Chen, Xianming Liu, and Junjun Jiang. DepthFormer: Exploiting long-range correlation and local information for accurate monocular depth estimation. *Machine Intelligence Research*, pages 1–18, 2023. 3
- [33] Zhaoshuo Li, Wei Ye, Dilin Wang, Francis X Creighton, Russell H Taylor, Ganesh Venkatesh, and Mathias Unberath. Temporally consistent online depth estimation in dynamic scenes. In *WACV*, 2023. 3
- [34] Zhenyu Li, Shariq Farooq Bhat, and Peter Wonka. PatchFusion: An end-to-end tile-based framework for high-resolution monocular metric depth estimation. In *CVPR*, 2024. 3
- [35] Zhenyu Li, Shariq Farooq Bhat, and Peter Wonka. PatchRefiner: Leveraging synthetic data for real-domain high-resolution monocular metric depth estimation. In *ECCV*, 2024. 3, 4
- [36] Ce Liu, Suryansh Kumar, Shuhang Gu, Radu Timofte, and Luc Van Gool. VA-DepthNet: A variational approach to single image depth prediction. In *ICLR*, 2023. 3
- [37] Shaoteng Liu, Yuechen Zhang, Wenbo Li, Zhe Lin, and Jiaya Jia. Video-P2P: Video editing with cross-attention control. In *CVPR*, 2024. 3
- [38] 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. *arXiv:2309.03453*, 2023. 4
- [39] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 5
- [40] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics*, 39(4), 2020. 3
- [41] 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
- [42] Jia Ning, Chen Li, Zheng Zhang, Chunyu Wang, Zigang Geng, Qi Dai, Kun He, and Han Hu. All in tokens: Unifying output space of visual tasks via soft token. In *ICCV*, 2023. 3
- [43] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. DINOv2: Learning robust visual features without supervision. *arXiv:2304.07193*, 2023. 2, 3
- [44] Emanuele Palazzolo, Jens Behley, Philipp Lottes, Philippe Giguere, and Cyrill Stachniss. ReFusion: 3d reconstruction in dynamic environments for RGB-D cameras exploiting residuals. In *IROS*, 2019. 6
- [45] Vaishakh Patil, Christos Sakaridis, Alexander Liniger, and Luc Van Gool. P3Depth: Monocular depth estimation with a piecewise planarity prior. In *CVPR*, 2022. 3
- [46] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. UniDepth: Universal monocular metric depth estimation. In *CVPR*, 2024. 3, 4
- [47] Luigi Piccinelli, Christos Sakaridis, Yung-Hsu Yang, Mattia Segu, Siyuan Li, Wim Abbeloos, and Luc Van Gool. UniDepthV2: Universal monocular metric depth estimation made simpler. *arXiv preprint arXiv:2502.20110*, 2025. 3
- [48] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. FateZero: Fusing attentions for zero-shot text-based video editing. In *ICCV*, 2023. 3
- [49] 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 TPAMI*, 2020. 2, 3, 4, 6
- [50] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M. Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *ICCV*, 2021. 5
- [51] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3
- [52] Saurabh Saxena, Abhishek Kar, Mohammad Norouzi, and David J Fleet. Monocular depth estimation using diffusion models. *arXiv:2302.14816*, 2023. 3
- [53] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 1
- [54] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. LAION-5B: An open large-scale dataset for training next generation image-text models. *NeurIPS*, 2022. 3
- [55] 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. 2, 3, 6
- [56] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 3
- [57] Zhonghao Sun, Wei Ye, Jinhui Xiong, Gyeongmin Choe, Jialiang Wang, Shuochen Su, and Rakesh Ranjan. Consistent direct time-of-flight video depth super-resolution. *arXiv:2211.08658*, 2022. 6
- [58] Zachary Teed and Jia Deng. Deepv2d: Video to depth with differentiable structure from motion. In *ICLR*, 2020. 3
- [59] Wenshan Wang, DeLong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. TartanAir: A dataset to push the limits of visual SLAM. In *IROS*, 2020. 5
- [60] Yiran Wang, Zhiyu Pan, Xingyi Li, Zhiguo Cao, Ke Xian, and Jianming Zhang. Less is more: Consistent video depth estimation with masked frames modeling. In *ACM MM*, 2022. 3
- [61] Yiran Wang, Min Shi, Jiaqi Li, Zihao Huang, Zhiguo Cao, Jianming Zhang, Ke Xian, and Guosheng Lin. Neural video depth stabilizer. In *ICCV*, 2023. 3, 6
- [62] Yiran Wang, Min Shi, Jiaqi Li, Chaoyi Hong, Zihao Huang, Juewen Peng, Zhiguo Cao, Jianming Zhang, Ke Xian, and

- Guosheng Lin. NVDS<sup>+</sup>: Towards efficient and versatile neural stabilizer for video depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 3, 6
- [63] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-A-Video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, 2023. 3
- [64] Guangkai Xu, Yongtao Ge, Mingyu Liu, Chengxiang Fan, Kangyang Xie, Zhiyue Zhao, Hao Chen, and Chunhua Shen. Diffusion models trained with large data are transferable visual models. *arXiv:2403.06090*, 2024. 3
- [65] Guanglei Yang, Hao Tang, Mingli Ding, Nicu Sebe, and Elisa Ricci. Transformer-based attention networks for continuous pixel-wise prediction. In *ICCV*, 2021. 3
- [66] Honghui Yang, Di Huang, Wei Yin, Chunhua Shen, Haifeng Liu, Xiaofei He, Binbin Lin, Wanli Ouyang, and Tong He. Depth any video with scalable synthetic data. *arXiv:2410.10815*, 2024. 3
- [67] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth Anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, 2024. 3, 6
- [68] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth Anything V2. *arXiv:2406.09414*, 2024. 2, 3, 4, 6
- [69] Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. Rerender a video: Zero-shot text-guided video-to-video translation. In *SIGGRAPH Asia*, 2023. 3
- [70] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. MVSNet: Depth inference for unstructured multi-view stereo. In *ECCV*, 2018. 2
- [71] Rajeev Yasarla, Hong Cai, Jisoo Jeong, Yunxiao Shi, Risheek Garrepalli, and Fatih Porikli. MAMo: Leveraging memory and attention for monocular video depth estimation. In *ICCV*, 2023. 3
- [72] Rajeev Yasarla, Manish Kumar Singh, Hong Cai, Yunxiao Shi, Jisoo Jeong, Yinhao Zhu, Shizhong Han, Risheek Garrepalli, and Fatih Porikli. FutureDepth: Learning to predict the future improves video depth estimation. In *European Conference on Computer Vision*, pages 440–458. Springer, 2024. 3
- [73] Wei Yin, Xinlong Wang, Chunhua Shen, Yifan Liu, Zhi Tian, Songcen Xu, Changming Sun, and Dou Renyin. DiverseDepth: Affine-invariant depth prediction using diverse data. *arXiv:2002.00569*, 2020. 3
- [74] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3D: Towards zero-shot metric 3d prediction from a single image. In *ICCV*, 2023. 3, 4
- [75] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. NeWCRFs: Neural window fully-connected CRFs for monocular depth estimation. In *CVPR*, 2022. 3
- [76] Haokui Zhang, Chunhua Shen, Ying Li, Yuanzhouhan Cao, Yu Liu, and Youliang Yan. Exploiting temporal consistency for real-time video depth estimation. In *ICCV*, 2019. 3
- [77] Xiang Zhang, Bingxin Ke, Hayko Riemenschneider, Nando Metzger, Anton Obukhov, Markus Gross, Konrad Schindler, and Christopher Schroers. BetterDepth: Plug-and-play diffusion refiner for zero-shot monocular depth estimation. *NeurIPS*, 2024. 3
- [78] Zhoutong Zhang, Forrester Cole, Richard Tucker, William T Freeman, and Tali Dekel. Consistent depth of moving objects in video. *ACM Transactions on Graphics (TOG)*, 40(4):1–12, 2021. 3
- [79] Zicheng Zhang, Bonan Li, Xuecheng Nie, Congying Han, Tiande Guo, and Luoqi Liu. Towards consistent video editing with text-to-image diffusion models. *NeurIPS*, 2024. 3
- [80] Wenliang Zhao, Yongming Rao, Zuyan Liu, Benlin Liu, Jie Zhou, and Jiwen Lu. Unleashing text-to-image diffusion models for visual perception. *arXiv:2303.02153*, 2023. 3
- [81] Zixiang Zhao, Jianshe Zhang, Shuang Xu, Zudi Lin, and Hanspeter Pfister. Discrete cosine transform network for guided depth map super-resolution. In *CVPR*, pages 5697–5707, 2022. 3
- [82] Zixiang Zhao, Haowen Bai, Yuanzhi Zhu, Jianshe Zhang, Shuang Xu, Yulun Zhang, Kai Zhang, Deyu Meng, Radu Timofte, and Luc Van Gool. DDFM: Denoising diffusion model for multi-modality image fusion. In *ICCV*, 2023. 3
- [83] 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 *ICCV*, 2023. 5