This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Continuous 3D Perception Model with Persistent State

Qianqian Wang^{1,2*}, Yifei Zhang^{1*}, Aleksander Holynski^{1,2}, Alexei A. Efros¹, Angjoo Kanazawa¹

¹University of California, Berkeley

²Google DeepMind



Static Scene

Dynamic Scene

Sparse Photo Collection

Figure 1. **Continuous 3D Perception.** Given a stream of RGB images as input, our approach enables dense 3D reconstruction in an online, continuous manner, estimating both camera parameters and dense 3D geometry with each incoming frame. Our framework supports various 3D tasks, processes inputs from video sequences and sparse photo collections, and can handle both static and dynamic scenes.

Abstract

We present a unified framework capable of solving a broad range of 3D tasks. Our approach features a stateful recurrent model that continuously updates its state representation with each new observation. Given a stream of images, this evolving state can be used to generate metric-scale pointmaps (per-pixel 3D points) for each new input in an online fashion. These pointmaps reside within a common coordinate system, and can be accumulated into a coherent, dense scene reconstruction that updates as new images arrive. Our model, called CUT3R (Continuous Updating Transformer for 3D) **R**econstruction), captures rich priors of real-world scenes: not only can it predict accurate pointmaps from image observations, but it can also infer unseen regions of the scene by probing at virtual, unobserved views. Our method is simple yet highly flexible, naturally accepting varying length of images that may be either video streams or unordered photo collections, containing both static and dynamic content. We evaluate our method on various 3D/4D tasks and demonstrate competitive or state-of-the-art performance in each. Project page: https://cut3r.github.io/.

1. Introduction

Humans are online visual learners. We continuously process streams of visual input, building on what we have learned in the past while learning in the present. Our prior knowledge enables us to interpret the world from minimal information; *e.g.*, upon entering a new restaurant, it only takes a glance to start inferring its layout and atmosphere. But it doesn't stop there—as we accumulate more observations, we continuously refine our mental model of the 3D environment. This ability to reconcile our prior knowledge of the world with a continuous stream of new observations is crucial for functioning effectively in an ever-changing visual world.

Building on these insights, we introduce an online 3D perception framework that unifies three key capabilities: 1) reconstructing 3D scenes from few observations, 2) continuously refining the reconstruction with more observations, and 3) inferring 3D properties of unobserved scene regions. We achieve these capabilities by integrating data-driven priors with a recurrent update mechanism. The learned prior enables our method to address challenges encountered by traditional methods (e.g., dynamic objects, sparse observations, degenerate camera motion), while the ability to continuously update allows it to process new observations online, and improve the reconstruction continuously over time.

Specifically, given an image stream, our recurrent model



Figure 2. **Querying Unseen Regions.** In addition to reconstructing a scene from images, our method can also infer structure for *unseen* parts of the scene, given a virtual camera query (shown in blue).

maintains and incrementally updates a persistent internal state that encodes the scene content. With each new observation, the model simultaneously updates this state and reads from it to predict the current view's 3D properties, including an estimate of that view's dense 3D geometry (as a pointmap; a 3D point per-pixel in a world coordinate frame) and camera parameters (both intrinsics and extrinsics). Accumulating these pointmaps enables online dense scene reconstruction, as illustrated in Fig. 1. Additionally, our framework supports inferring unobserved parts of the scene: by querying the internal state with a virtual (unseen) view, parameterized as a raymap, we can extract the corresponding pointmap and color for the query view, as depicted in Fig. 2.

Our framework is designed to be general and flexible, making it well-suited for training on an extensive collection of datasets and adaptable to diverse inference scenarios. During training, we leverage a wide variety of 3D data, including single images, videos, and photo collections with partial or full 3D annotations. These datasets span a broad spectrum of scene types and contexts-static and dynamic, indoor and outdoor, real and synthetic-enabling the model to acquire robust and generalizable priors. During inference, our recurrent framework naturally accepts varying numbers of images, and supports a wide range of input data settings: from streaming video to unstructured image collections, including wide-baseline or even non-overlapping images. Bevond static scenes, it seamlessly handles videos of dynamic scenes, estimating accurate camera parameters and dense point clouds for moving parts of the scene.

We evaluate our method on various 3D tasks: monocular and consistent video depth estimation, camera pose estimation, and 3D reconstruction, achieving competitive or state-of-the-art performance in each. We also show that our method can infer previously unseen structures and continuously refine the reconstruction as new observations arrive.

2. Related Work

Tabula rasa **3D Reconstruction.** Many 3D reconstruction pipelines operate as *tabula rasa*¹, starting from scratch for

each new scene, and relying solely on the observations currently available. Notable examples include traditional methods such as Structure from Motion (SfM) [1, 2, 28, 47, 65, 66, 70, 71, 80] and Simultaneous Localization and Mapping (SLAM) [11, 16, 19, 22, 37, 51, 53], as well as more recent approaches such as Neural Radiance Fields (NeRF) [12, 24, 49, 50, 84] and 3D Gaussian Splatting [36]. The *tabula rasa* nature of these approaches poses challenges in handling scenarios with sparse observations or ill-posed conditions. In contrast, our method leverages data-driven priors to enable dense 3D reconstruction directly from video sequences or photo collections, eliminating the need for known camera extrinsics or intrinsics. The data-driven nature of our method enables it to address challenging cases of degeneracy, make predictions from as few as a single image, and infer structures that are unobserved in the input.

Learning-Based 3D Reconstruction. Unlike tabula rasa reconstruction, many methods integrate data-driven priors into 3D reconstruction. One prominent direction focuses on improving traditional reconstruction pipelines, including replacing hand-crafted components with learning-based alternatives (e.g., substituting conventional feature descriptors [6, 47] with learned ones [17, 20, 63, 73]), integrating data-driven priors [13, 77-79, 93, 97, 105] into the systems, or optimizing entire systems end-to-end [76, 78, 83, 94]. Another substantial body of work aims to predict dense 3D geometry directly from single or paired images. Extensive research [7, 27, 42, 55, 58, 59, 85, 92] focuses on estimating monocular depth. While depth maps provide valuable 3D information, converting them into 3D point clouds typically requires camera intrinsics [8, 32, 55], with camera extrinsics often left unmodeled. Notably, DUSt3R [86] predicts two pointmaps from an image pair within the same coordinate frame, inherently accounting for both camera intrinsics and extrinsics. However, DUSt3R is tailored for image pairs and does not inherently support multiple views. While it can be extended to handle multi-view tasks via an additional global alignment process, this process operates offline, is time-consuming, and cannot dynamically update the reconstruction as new observations are added. In contrast, our method flexibly handles a varying number of images, predicting pointmaps for each image in a shared coordinate frame as they are received.

Continuous Reconstruction Methods. Many traditional and learning-based methods share the ability to continuously predict 3D structure in an online manner. Monocular SLAM pipelines [22, 23, 78, 105] recover ego-motion and 3D point clouds in real-time from video sequences but typically require known camera intrinsics. Our approach is related to learning based methods [14, 34, 74, 100], such as 3D-R2N2 [14], which utilize recurrent neural network architectures [21, 31, 33] for online 3D reconstruction. However,

¹tabula rasa (lat.) – blank slate.

these methods are either object-centric [14, 34, 96] or require posed images as input [9, 34, 64, 74, 100]. Spann3R [81], developed concurrently with our work, also demonstrates continuous reconstruction capabilities using a spatial memory mechanism. However, while Spann3R's memory serves primarily as a cache for observed scenes, our compressed state representation not only captures observed scene content but also enables inferring unobserved structures. Finally, unlike the methods discussed in this paragraph that only operate on static scenes, our method also seamlessly reconstructs dynamic scenes.

Reconstructing Dynamic Scenes from Monocular Videos. Recovering camera parameters and consistent dense geometry from monocular videos of dynamic scenes presents significant challenges for traditional SLAM pipelines. Recent approaches address this by leveraging learned depth priors. Robust-CVD [38] uses deformation splines to align the depth maps, while CasualSAM [101] finetunes a monocular depth network on a single video. These methods involve time-consuming per-video optimization. Concurrent to our work, MonST3R [98] extends DUSt3R to predict pointmaps for dynamic scenes by finetuning it on dynamic datasets. However, it still follows DUSt3R's pairwise formulation and requires global alignment as post-processing. Another concurrent work, MegaSaM [43], achieves highly accurate and robust estimation of camera poses and scene structure for casually captured dynamic videos. Unlike our approach, MegaSaM is optimization-based (i.e., not feedforward), uses an explicit 3D state and does not make online predictions.

3D Scene Priors. Predicting 3D content beyond observed views has long been studied, but it has recently gained significant attention due to advancements in generative modeling. Regression-based few-shot novel view synthesis approaches [39, 68, 69, 96] typically only generalize across a class of 3D scenes. With rapid advances in image and video generative models [57, 61, 62], much of the current research in 3D generation focuses on transferring priors from images and videos to 3D [45, 48, 56, 91]. However, most imagebased 3D generation methods require camera parameters as input, focus on limited data domains (e.g., object-centric), or require additional 3D distillation processes to extract 3D content. In contrast to these view-centric approaches that generate novel views as a proxy for 3D, our method is geometrycentric: It directly generates metric-scale pointmaps for a virtual camera query in a scene reconstructed from a set of images, without requiring their camera intrinsics or poses.

3. Method

Our approach takes as input a stream of images without any camera information. The image streams can come from either video or image collections. As a new image comes in through the model, it interacts with the latent state representation, which encodes the understanding of the current 3D scene. Specifically, the image simultaneously updates the state with new information and retrieves information stored in the state. Following the state-image interaction, explicit 3D pointmaps and camera poses are extracted for each view. The state can also be queried with a virtual view to predict its corresponding pointmap, capturing unseen parts of the scene. See Fig. 3 for our method overview.

3.1. State-Input Interaction Mechanism

Our method takes a stream of images as input. For each current image, I_t , it is first encoded into token representation by a ViT encoder [18]:

$$\boldsymbol{F}_t = \text{Encoder}_i(\boldsymbol{I}_t). \tag{1}$$

We represent the state also as a set of tokens. Prior to seeing any image input, the state tokens are initialized as a set of learnable tokens shared by all scenes. The image tokens interact with the state in two ways: they update the state with information from the current image and read the context from the state, incorporating stored past information. We refer to these interactions as *state-update* and *state-readout*, respectively. This bidirectional interaction is implemented using two interconnected transformer decoders [86, 88], which jointly operate on both image and state tokens:

$$[\boldsymbol{z}_t', \boldsymbol{F}_t'], \boldsymbol{s}_t = \text{Decoders}([\boldsymbol{z}, \boldsymbol{F}_t], \boldsymbol{s}_{t-1}). \tag{2}$$

Here s_{t-1} and s_t represent the state tokens before and after interaction with the image tokens. F'_t denotes the image tokens enriched with state information. z is a learnable "pose token" prepended to the image tokens, whose output z'_t captures image-level information related to the scene, such as ego motion. Within the decoders, the outputs from both sides cross-attend to each other at each decoder block to ensure effective information transfer.

After this interaction, explicit 3D representation can be extracted from F'_t and z'_t . Specifically, we predict two pointmaps with corresponding confidence maps, $(\hat{X}^{\text{self}}_t, C^{\text{self}}_t)$ and $(\hat{X}^{\text{world}}_t, C^{\text{world}}_t)$. These maps are defined in two coordinate frames: the input image's own coordinate frame and the world frame, respectively, where the world frame is defined as the coordinate frame of the initial image. Additionally, we predict the relative transformations between the two coordinate frames, or, the ego motion, \hat{P}_t :

$$\hat{X}_{t}^{\text{self}}, C_{t}^{\text{self}} = \text{Head}_{\text{self}}(F_{t}')$$
 (3)

$$\hat{X}_{t}^{\text{world}}, C_{t}^{\text{world}} = \text{Head}_{\text{world}}(F_{t}', z_{t}')$$
 (4)

$$\hat{P}_t = \text{Head}_{\text{pose}}(\boldsymbol{z}'_t),$$
 (5)

where $Head_{self}$ and $Head_{world}$ are implemented as DPT [58], and $Head_{pose}$ is implemented as an MLP network, respectively. Please see the supplement for details. We extract



Figure 3. **Method Overview**. Our method performs online dense 3D reconstruction from a stream of images (video frames or a photo collection) by using a persistent state. Each input image is encoded into visual tokens via a shared-weight ViT encoder. These tokens interact with state tokens, where *state update* integrates the current image into the state, and *state readout* retrieves the past context stored in the state for predictions. Both processes occur simultaneously through two interconnected ViT decoders. Outputs include pointmaps in world and camera frames (only world pointmaps are shown) and the camera-to-world transformation. On the right, we demonstrate our method's ability to predict unseen views: given a query camera (as a raymap), it reads information from the state to predict its corresponding pointmap, even for unobserved regions. For these readouts, we do not update the state. The hallucinated pointmap is highlighted with a blue border.

6-DoF pose \hat{P}_t from the pose token z'_t , the rigid transformation from the current frame to the world. All pointmaps and poses are in metric scale (i.e., meters).

Although predicting \hat{X}_t^{self} , \hat{X}_t^{world} , and \hat{P}_t may seem redundant, we found this redundancy simplifies training. It enables each output to receive direct supervision, and importantly, it facilitates training on datasets with partial annotations, such as those containing only pose or single-view depth, thereby broadening the range of usable data.

3.2. Querying the State with Unseen Views

Leveraging past 3D experiences, humans can envision parts of a scene beyond what is directly observed. We emulate this ability by extending the state-readout operation to predict unseen portions of the scene from a virtual camera view. Specifically, we use a virtual camera as a query to extract information from the state. The virtual camera's intrinsics and extrinsics are represented as a raymap R, a 6-channel image encoding the origin and direction of rays at each pixel [25, 90, 99].

Given a query raymap R, we first encode it into token representations F_r using a separate transformer Encoder_r:

$$\boldsymbol{F}_r = \operatorname{Encoder}_r(\boldsymbol{R}).$$
 (6)

Then, the rest of the process aligns largely with what is described in Sec 3.1. Specifically, we have F_r interact with the current state through the same decoder module (i.e., shared weights) as in Eq. 2 to read from the state into F'_r . Note that, unlike in the state-image interaction, the state is not

updated here, as the raymap serves solely as a query without introducing new scene content. Finally, we apply the same head networks to parse F'_r into explicit representations, as in Eq. 3-5. Additionally, we introduce another head Head_{color} to decode color information: $\hat{I}_r = \text{Head}_{\text{color}}(F'_r)$ which corresponds to the color of each ray in the raymap R.

Querying the scene with raymaps has an interesting analogy to Masked Autoencoders (MAE) [30]. In MAE, completion occurs at the patch level, using the global context of the entire image. Here, completion is performed at the image level, leveraging the global context of the 3D scene captured in the state.

3.3. Training Objective

During training, we provide the model with a sequence of N images (either from a video sequence or an image collection). The raymap mode is enabled only for training data with metric-scale 3D annotations. In these cases, we randomly replace each image with its corresponding raymap at a certain probability, excluding the first view. When the scale of the 3D annotation is unknown, raymap querying is disabled to avoid inconsistencies between the annotation scale and the scene scale represented in the state. For clarity, we do not differentiate between predictions generated from a raymap or an image in this section to simplify the notation. We denote the pointmap predictions of the current training sequence as $\mathcal{X} = \{\hat{\mathcal{X}}^{\text{self}}, \hat{\mathcal{X}}^{\text{world}}\}$, where $\hat{\mathcal{X}}^{\text{self}} = \{\hat{\mathcal{X}}^{\text{self}}_t\}_{t=1}^N$, $\hat{\mathcal{X}}^{\text{world}} = \{\hat{\mathcal{X}}^{\text{self}}_t\}_{t=1}^N$ and their corresponding confidence scores as \mathcal{C} .

3D regression loss. Following MASt3R [40], we apply a confidence-aware regression loss to the pointmaps:

$$\mathcal{L}_{\text{conf}} = \sum_{(\hat{\boldsymbol{x}}, c) \in (\hat{\mathcal{X}}, \mathcal{C})} \left(c \cdot \left\| \frac{\hat{\boldsymbol{x}}}{\hat{s}} - \frac{\boldsymbol{x}}{s} \right\|_2 - \alpha \log c \right), \quad (7)$$

where \hat{s} and s are scale normalization factors for \hat{X} and \hat{X} , respectively. Similar to MASt3R [40], when the groundtruth pointmaps are metric, we set $\hat{s} := s$ to enable the model to learn metric-scale pointmaps.

Pose loss. We parameterize the pose \hat{P}_t as quaternion \hat{q}_t and translation $\hat{\tau}_t$, and minimize the L2 norm between the prediction and groundtruth:

$$\mathcal{L}_{\text{pose}} = \sum_{t=1}^{N} \left(\|\hat{\boldsymbol{q}}_t - \boldsymbol{q}_t\|_2 + \left\| \frac{\hat{\boldsymbol{\tau}}_t}{\hat{s}} - \frac{\boldsymbol{\tau}_t}{s} \right\|_2 \right).$$
(8)

RGB loss. When the input is raymap, besides the 3D regression loss, we also apply an MSE loss to enforce the predicted pixel colors \hat{I}_r to match the groundtruth: $\mathcal{L}_{rgb} = ||\hat{I}_r - I_r||_2^2$.

3.4. Training Strategy

Training Datasets. We train our method on a diverse set of 32 datasets, covering synthetic and real-world data, static and dynamic scenes, scene-level and object-centric configurations, as well as both indoor and outdoor scenes. Examples of our datasets include CO3Dv2 [60], ARK-itScenes [5], ScanNet++ [95], TartanAir [87], Waymo [75], MegaDepth [41], MapFree [3], DL3DV [44], and DynamicStereo [35]. Our flexible formulation allows training on datasets with partial annotations (*i.e.* only camera parameters like RealEstate10K [103], or only single views, like Synscapes [89]). See the supplement for the full list.

Curriculum Training. Our model is trained with a curriculum. The first stage trains the model on 4-view sequences from mainly static datasets. The second stage incorporates dynamic scene datasets, improving the model's ability to handle moving objects such as humans, along with datasets with partial annotations, further enhancing its generalization. These two stages are trained on 224×224 images to reduce computational costs, following DUSt3R [86]. In the third stage, we train with higher resolution, using varied aspect ratios and setting the maximum side to 512 pixels. Finally, we freeze the encoder, training only the decoder and heads on longer sequences spanning 4 to 64 views. This stage focuses on enhancing inter-scene reasoning and effectively handling long contexts.

Implementation Detail. We use a ViT-Large model [18] for the image encoder $Encoder_i$, initialized with DUSt3R encoder pretrained weights, and ViT-Base for the decoders. Both the encoder and decoders operate on 16×16 pixel patches. The state consists of 768 tokens, each with a dimensionality of 768. The raymap encoder $Encoder_r$ is a lightweight encoder with 2 blocks. We use Adam-W optimizer [46] with an initial learning rate of $1e^{-4}$, applying linear warmup followed by cosine decay. We train our model on eight A100 NVIDIA GPUs each with 80G memory. Please refer to the supplement for more details.

4. Experiments

We evaluate our method across a range of 3D tasks, including single and video depth estimation (Sec. 4.1), camera pose estimation (Sec. 4.2), and 3D reconstruction (Sec. 4.3).

Baselines. Our primary set of baselines are DUSt3R [86], MASt3R [40], Spann3R [81], and MonST3R [98], where the latter two are concurrent works. MonST3R finetunes DUSt3R on dynamic datasets to handle dynamic scenes, while Spann3R extends DUSt3R to support varying number of images via additional spatial memory and operates online, similar to our method. DUSt3R, MASt3R, and MonST3R can only take a pair of views as input, and require an extra global alignment (GA) stage to consolidate the pairwise predictions. Both MASt3R and our method predict metric pointmaps, whereas others predict relative pointmaps.

4.1. Monocular and Video Depth Estimation

Mono-Depth Estimation. Following MonST3R [98], we evaluate monocular depth estimation on KITTI [26], Sintel [10], Bonn [54] and NYU-v2 [52] datasets covering dynamic and static, indoor and outdoor scenes. These datasets are excluded from training, enabling zero-shot performance evaluation across domains. We use absolute relative error (Abs Rel) and $\delta < 1.25$ (percentage of predicted depths within a 1.25-factor of true depth) as metrics, with per-frame median scaling per DUSt3R [86]. Results in Tab. 1 show our method achieves state-of-the-art or competitive performance, leading on Bonn and and NYU-v2 and ranking second on KITTI.

	Sintel		Bo	nn	Kľ	ГТІ	NYU-v2		
Method	Abs Rel↓	$\delta < 1.25 \uparrow$	Abs Rel↓	$\delta {<} 1.25 \uparrow$	Abs Rel↓	$\delta {<} 1.25 \uparrow$	Abs Rel↓	δ<1.25	
DUSt3R	0.424	58.7	0.141	82.5	0.112	86.3	0.080	90.7	
MASt3R	0.340	60.4	0.142	82.0	0.079	94.7	0.129	84.9	
MonST3R	0.358	54.8	0.076	93.9	0.100	89.3	0.102	88.0	
Spann3R	0.470	53.9	0.118	85.9	0.128	84.6	0.122	84.9	
Ours	0.428	55.4	0.063	96.2	0.092	<u>91.3</u>	0.086	90.9	

Table 1. **Single-frame Depth Evaluation.** We report the performance on Sintel, Bonn, KITTI, and NYU-v2 (static) datasets.

Video Depth Estimation. Video depth estimation evaluates per-frame depth quality and inter-frame depth consistency by aligning predicted depth maps to ground truth using a per-sequence scale. For metric pointmap methods like ours and MASt3R, we also report results without alignment. Comparisons for both methods are presented in Tab. 2.

Under per-sequence scale alignment, our method consistently outperforms DUSt3R [86] and MASt3R [40]. The

			Sir	ntel	BO	NN	КП		
Alignment	Method	Optim. Onl.	Abs Rel \downarrow	$\delta{<}1.25\uparrow$	Abs Rel↓	$\delta{<}1.25\uparrow$	Abs Rel \downarrow	$\delta < 1.25 \uparrow$	FPS
	DUSt3R-GA [86]	\checkmark	0.656	45.2	0.155	83.3	0.144	81.3	0.76
	MASt3R-GA [40]	\checkmark	0.641	43.9	0.252	70.1	0.183	74.5	0.31
D I	MonST3R-GA [98]	\checkmark	0.378	55.8	0.067	96.3	0.168	74.4	0.35
Per-sequence scale	Spann3R [81]	\checkmark	0.622	42.6	0.144	81.3	0.198	73.7	13.55
	Ours	\checkmark	0.421	<u>47.9</u>	0.078	<u>93.7</u>	0.118	88.1	16.58
	MASt3R-GA [40]	\checkmark	1.022	14.3	0.272	70.6	0.467	15.2	0.31
Metric scale	Ours	\checkmark	1.029	23.8	0.103	88.5	0.122	85.5	16.58

Table 2. Video Depth Evaluation. We report scale-invariant depth and metric depth accuracy on Sintel, Bonn, and KITTI datasets. Methods requiring global alignment are marked "GA", while "Optim." and "Onl." indicate optimization-based and online methods, respectively. We also report the FPS on KITTI dataset using 512×144 image resolution for all methods on an A100 GPU, except Spann3R which only supports 224×224 inputs. We present a subset of baselines here; please refer to the supplementary material for full comparisons.

			Sintel				TUM-dynan	nics	ScanNet			
Method	Optim.	Onl.	$ATE\downarrow$	RPE trans \downarrow	RPE rot \downarrow	$ $ ATE \downarrow	RPE trans \downarrow	RPE rot \downarrow	$ $ ATE \downarrow	RPE trans \downarrow	RPE rot ↓	
Particle-SfM [102]	\checkmark		0.129	0.031	0.535	-	-	-	0.136	0.023	0.836	
Robust-CVD [38]	\checkmark		0.360	0.154	3.443	0.153	0.026	3.528	0.227	0.064	7.374	
CasualSAM [101]	\checkmark		0.141	0.035	0.615	<u>0.071</u>	0.010	1.712	0.158	0.034	1.618	
DUSt3R-GA [86]	\checkmark		0.417	0.250	5.796	0.083	0.017	3.567	0.081	0.028	0.784	
MASt3R-GA [40]	\checkmark		0.185	0.060	1.496	0.038	0.012	0.448	0.078	0.020	0.475	
MonST3R-GA [98]	\checkmark		0.111	0.044	0.869	0.098	0.019	0.935	0.077	0.018	0.529	
DUSt3R [86]		\checkmark	0.290	0.132	7.869	0.140	0.106	3.286	0.246	0.108	8.210	
Spann3R [81]		\checkmark	0.329	0.110	4.471	<u>0.056</u>	0.021	0.591	0.096	0.023	0.661	
Ours		\checkmark	0.213	0.066	0.621	0.046	0.015	0.473	<u>0.099</u>	0.022	0.600	

Table 3. Evaluation on Camera Pose Estimation on Sintel [10], TUM-dynamic [72], and ScanNet [15] datasets. Our method achieves the best overall performance among all online methods.

global alignment they use assumes that the scene is static, and enforcing multi-view consistency can only improve the reconstruction of static regions but may impair the reconstruction of moving objects. In contrast, our method leverages implicit alignment through the state, making it adaptable to both static and dynamic scenes while remaining fully online. Our method also significantly outperforms the other online method, Spann3R [81], which is based on spatial memory designed with a static scene assumption and trained only on static 3D scenes. MonST3R [98] achieves stateof-the-art performance, but it relies on an additional global alignment stage that incorporates extra input like optical flow in its optimization. In comparison, our method performs comparably, or even better on the KITTI dataset, while remaining online and achieving nearly $50 \times$ speedup. In the metric-scale setting, our method also significantly outperforms MASt3R for most metrics.

4.2. Camera Pose Estimation

Following MonST3R [98], we evaluate camera pose estimation accuracy on Sintel [10], TUM dynamics [72], and Scan-Net [15] datasets. Note that both Sintel and TUM-dynamics contain major dynamic objects, making it challenging for traditional SfM and SLAM systems. We report Absolute Translation Error (ATE), Relative Translation Error (RPE trans), and Relative Rotation Error (RPE rot) after Sim(3) alignment with the ground truth, as in [13, 98, 102]. Unlike most visual odometry methods [13, 29, 78], our method does not require any camera calibration. We compare to baselines that share this feature. Most prior approaches do so through test-time optimization, as seen in Robust-CVD [38] and CasualSAM [101], which jointly estimate camera parameters and dense depth maps per sequence.

The results are presented in Tab. 3. We separately highlight the leading approaches for methods that require additional optimization and those that do not (i.e., online). For the online category, we additionally include DUSt3R [86] where we align all video frames with first frame, without using GA. Although a gap persists between optimizationbased and online methods, our approach achieves the best overall performance among online methods, particularly in dynamic scenes.

4.3. 3D Reconstruction

We evaluate scene-level reconstruction on the 7-scenes [67] and NRGBD [4] datasets using accuracy (Acc), completion (Comp), and normal consistency (NC) metrics, as in prior works [4, 81, 82, 86, 104]. To assess performance on image collections with minimal or no overlap, we evaluate using sparsely sampled images: 3 to 5 frames per scene for the 7-Scenes dataset and 2 to 4 frames per scene for the NRGBD dataset. The results are presented in Tab. 4. Our



Figure 4. **Qualitative Results on In-the-wild Internet Videos.** We compare our method with concurrent works Spann3R [81] and MonST3R [98]. Our method achieves the best qualitative results.

			7 scenes [67]						NRGBD [4]						
			Ac	c↓	Cor	np↓	N	Z↑	Ac	c↓	Cor	np↓	N		
Method	Optim.	Onl.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	FPS
DUSt3R-GA [86]	\checkmark		0.146	0.077	0.181	0.067	0.736	0.839	0.144	0.019	0.154	0.018	0.870	0.982	0.68
MASt3R-GA [40]	\checkmark		0.185	0.081	0.180	0.069	0.701	0.792	0.085	0.033	0.063	0.028	0.794	0.928	0.34
MonST3R-GA [40]	\checkmark		0.248	0.185	0.266	0.167	0.672	0.759	0.272	0.114	0.287	0.110	0.758	0.843	0.39
Spann3R [81]		\checkmark	0.298	0.226	0.205	0.112	0.650	0.730	0.416	0.323	0.417	0.285	0.684	0.789	12.97
Ours		\checkmark	0.126	0.047	0.154	0.031	0.727	<u>0.834</u>	0.099	<u>0.031</u>	<u>0.076</u>	<u>0.026</u>	<u>0.837</u>	<u>0.971</u>	17.00

Table 4. **3D reconstruction comparison on 7-Scenes [67] and NRGBD [4] datasets.** While operating online, our method achieves competitive performance, on par with and even surpassing offline methods that employ global alignment.

method significantly outperforms the other online approach Spann3R [81], and achieves comparable or sometimes better results than the top optimization-based method, DUSt3R-GA, while operating online at $25 \times$ the speed. This highlights our method's effectiveness with sparse image collections.

Qualitative Results. We compare the reconstruction quality of our method with Spann3R [81] and MonST3R [98] on in the wild Internet videos in Fig. 4. Spann3R [81] is neither designed nor trained on dynamic scenes, making it less effective at handling moving objects, such as humans. MonST3R [98] is finetuned on dynamic scenes, potentially overfitting and degrading performance on static 3D scenes. In contrast, our method operates online and achieves state-of-the-art performance across both static and dynamic scenes.

4.4. Analysis

State Update Analysis. Our model continuously updates its state representation as new data arrives, relying solely on

past and current observations without knowledge of future inputs. As more observations accumulate, the state should be able to refine its understanding of the 3D world, leading to improved predictions. We demonstrate this capability of our method in Tab. 5. Using the same experimental setup as in Sec. 4.3, we introduce an additional version of our approach called "revisiting": we first run our method online to obtain the final state that has seen all images, then we freeze this state and use it to process the same set of images again to generate predictions.

This setup differs from the online setup by allowing the state to see the full context of the scene during the first run. As shown in Tab. 5, revisiting improves performance compared to the online version, especially for accuracy. This verifies that the state representation effectively updates with additional observations. See Fig. 5 for a qualitative example.

		7-Scenes		NRGBD					
Method	Acc↓	Comp↓	NC↑	Acc↓	Comp↓	NC↑			
DUSt3R-GA [86]	0.146	0.181	0.736	0.144	0.154	0.870			
Ours	0.126	0.154	0.727	0.099	0.076	0.837			
Ours Revisit	0.113	0.107	<u>0.732</u>	0.094	0.076	0.844			

Table 5. State Update Analysis on 7-Scenes [67] and NRGBD [4] datasets.



Figure 5. **State Update Analysis.** Compared to online, revisiting incorporates global context which improves overall reconstruction results, especially in the highlighted regions.

Inferring Unseen Regions via State Readout. To the best of our knowledge, our method is the first to enable the inference of unseen structures in metric scale for general scenes, supporting both single and multiple views, without requiring camera intrinsics or poses for the input images. We show qualitative results of the generated structures in Fig. 6. For this experiment, we use the validation set of the MapFree [3] and ARKitScenes datasets, both with metric camera pose annotations. Importantly, these scenes are not seen by our model during training. In each example, we input a single image to our model and then query the state using a raymap of the ground truth image (unseen by the state). The model is expected to generate pointmaps that align with the ground truth. While the predictions may lack some high-frequency details - owing to the deterministic nature of our approach they accurately follows the viewpoint transformation, even with significant viewpoint changes between the input and ground truth. In addition, our method generates new structures beyond what what is observed in the input, such as the bushes in the first example, the ground in the second, the oven in the third, and the stool in the last. This demonstrates that our method captures generalized 3D scene priors.

5. Conclusion

In this paper, we propose an online 3D perception model with a continuously updating, persistent state. Given an im-



Figure 6. **Inferring New Structure via. State Readout**. From top to bottom: the input image; the ground truth (GT) image, used to query the state via its camera parameters (note: GT image is not given to the model); the depth map from the predicted pointmap; the pointmap prediction of the input image alone; and the pointmap combined with the predicted pointmap in a shared coordinate frame.

age stream, our model simultaneously performs state-update (which updates the state) and state-readout (which retrieves information from the state) for each observation. The output at each step includes camera parameters and pointmaps in the world frame, which accumulate into a dense reconstruction of the scene over time. This simple formulation is general yet powerful enough to solve a number of 3D/4D tasks, handling both videos and photo collections, and processing both static and dynamic scenes. The generalized 3D scene priors captured by our method enable the inference of new structures unobserved in the input views by probing the state with a raymap. Experimental results on extensive 3D/4D tasks verify the effectiveness of our method.

Limitations. As with many online methods, our approach may eventually drift over very long sequences due to the absence of global alignment. Extending our work with explicit or implicit global alignment is an interesting future direction. Additionally, since our structure generation is performed via a deterministic rather than a generative approach, it can produce blurry results – especially when extrapolating viewpoints too far from the provided views – a common issue with regression-based methods. Incorporating a generative formulation could address this limitation. Finally, training recurrent networks can be time-consuming. We leave these fundamental directions for future work.

Acknowledgements. We would like to thank Noah Snavely, Haiwen Feng, Chung Min Kim, Justin Kerr, Songwei Ge, Chenfeng Xu, Letian Fu, and Ren Wang for helpful discussions. We especially thank Noah Snavely for his guidance and support. This project is supported in part by DARPA No. HR001123C0021, IARPA DOI/IBC No. 140D0423C0035, NSF:CNS-2235013, Bakar Fellows, ONR, MURI, TRI and BAIR sponsors. The views and conclusions contained herein are those of the authors and do not represent the official policies or endorsements of these institutions.

References

- Sameer Agarwal, Noah Snavely, Steven M Seitz, and Richard Szeliski. Bundle adjustment in the large. In Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part II 11, pages 29–42. Springer, 2010.
- [2] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54 (10):105–112, 2011. 2
- [3] Eduardo Arnold, Jamie Wynn, Sara Vicente, Guillermo Garcia-Hernando, Áron Monszpart, Victor Adrian Prisacariu, Daniyar Turmukhambetov, and Eric Brachmann. Map-free visual relocalization: Metric pose relative to a single image. In ECCV, 2022. 5, 8
- [4] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural rgb-d surface reconstruction. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 6290– 6301, 2022. 6, 7, 8
- [5] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, and Elad Shulman. ARKitscenes - a diverse real-world dataset for 3d indoor scene understanding using mobile RGB-d data. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. 5
- [6] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer* vision and image understanding, 110(3):346–359, 2008. 2
- [7] Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023. 2
- [8] 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, 2024. 2
- [9] Aljaz Bozic, Pablo Palafox, Justus Thies, Angela Dai, and Matthias Nießner. Transformerfusion: Monocular rgb scene reconstruction using transformers. Advances in Neural Information Processing Systems, 34:1403–1414, 2021. 3

- [10] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conf. on Computer Vision (ECCV)*, pages 611– 625. Springer-Verlag, 2012. 5, 6
- [11] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics*, 32(6):1309–1332, 2016. 2
- [12] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, pages 333–350. Springer, 2022. 2
- [13] Weirong Chen, Le Chen, Rui Wang, and Marc Pollefeys. Leap-vo: Long-term effective any point tracking for visual odometry. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 19844– 19853, 2024. 2, 6
- [14] 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 *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII 14*, pages 628–644. Springer, 2016. 2, 3
- [15] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richlyannotated 3d reconstructions of indoor scenes. In Proc. Computer Vision and Pattern Recognition (CVPR), IEEE, 2017. 6
- [16] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *IEEE transactions on pattern analysis and machine intelli*gence, 29(6):1052–1067, 2007. 2
- [17] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. 2
- [18] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020. 3, 5
- [19] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006. 2
- [20] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint description and detection of local features. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition*, pages 8092–8101, 2019. 2
- [21] Jeffrey L Elman. Finding structure in time. Cognitive science, 14(2):179–211, 1990. 2
- [22] Jakob Engel, Thomas Schöps, and Daniel Cremers. Lsdslam: Large-scale direct monocular slam. In *European conference on computer vision*, pages 834–849. Springer, 2014. 2
- [23] Christian Forster, Zichao Zhang, Michael Gassner, Manuel Werlberger, and Davide Scaramuzza. Svo: Semidirect visual

odometry for monocular and multicamera systems. *IEEE Transactions on Robotics*, 33(2):249–265, 2016. 2

- [24] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5501–5510, 2022. 2
- [25] 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. arXiv preprint arXiv:2405.10314, 2024. 4
- [26] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231– 1237, 2013. 5
- [27] Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3828–3838, 2019. 2
- [28] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 2
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016. 6
- [30] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 4
- [31] S Hochreiter. Long short-term memory. *Neural Computation MIT-Press*, 1997. 2
- [32] 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 preprint arXiv:2404.15506, 2024.
- [33] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021.
- [34] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. Advances in neural information processing systems, 30, 2017. 2, 3
- [35] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Dynamicstereo: Consistent dynamic depth from stereo videos. *CVPR*, 2023. 5
- [36] 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

- [37] Georg Klein and David Murray. Parallel tracking and mapping for small ar workspaces. In 2007 6th IEEE and ACM international symposium on mixed and augmented reality, pages 225–234. IEEE, 2007. 2
- [38] 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, pages 1611–1621, 2021. 3, 6
- [39] Zihang Lai, Sifei Liu, Alexei A Efros, and Xiaolong Wang. Video autoencoder: self-supervised disentanglement of static 3d structure and motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9730–9740, 2021. 3
- [40] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. arXiv preprint arXiv:2406.09756, 2024. 5, 6, 7
- [41] Zhengqi Li and Noah Snavely. Megadepth: Learning singleview depth prediction from internet photos. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [42] Zhengqi Li and Noah Snavely. Megadepth: Learning singleview depth prediction from internet photos. In *Proceedings* of the IEEE conference on computer vision and pattern recognition, pages 2041–2050, 2018. 2
- [43] 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. arXiv preprint arXiv:2412.04463, 2024. 3
- [44] Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. Dl3dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22160–22169, 2024. 5
- [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, 2023. 3
- [46] I Loshchilov. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101, 2017. 5
- [47] David G Lowe. Distinctive image features from scaleinvariant keypoints. *International journal of computer vision*, 60:91–110, 2004. 2
- [48] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Realfusion: 360 reconstruction of any object from a single image. In CVPR, 2023. 3
- [49] 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.
- [50] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. ACM transactions on graphics (TOG), 41(4):1–15, 2022. 2
- [51] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015. 2

- [52] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 5
- [53] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. Dtam: Dense tracking and mapping in real-time. In 2011 international conference on computer vision, pages 2320–2327. IEEE, 2011. 2
- [54] E. Palazzolo, J. Behley, P. Lottes, P. Giguère, and C. Stachniss. ReFusion: 3D Reconstruction in Dynamic Environments for RGB-D Cameras Exploiting Residuals. 2019. 5
- [55] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10106–10116, 2024. 2
- [56] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988, 2022. 3
- [57] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr, 2021. 3
- [58] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *ICCV*, 2021. 2, 3
- [59] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot crossdataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3), 2022. 2
- [60] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *International Conference on Computer Vision*, 2021. 5
- [61] 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, pages 10684–10695, 2022. 3
- [62] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 3
- [63] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In CVPR, 2020. 2
- [64] Mohamed Sayed, John Gibson, Jamie Watson, Victor Prisacariu, Michael Firman, and Clément Godard. Simplerecon: 3d reconstruction without 3d convolutions. In *European Conference on Computer Vision*, pages 1–19. Springer, 2022. 3

- [65] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [66] Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In 2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06), pages 519–528. IEEE, 2006. 2
- [67] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in rgb-d images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2930–2937, 2013. 6, 7, 8
- [68] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structureaware neural scene representations. Advances in Neural Information Processing Systems, 32, 2019. 3
- [69] Vincent Sitzmann, Semon Rezchikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems*, 34: 19313–19325, 2021. 3
- [70] Noah Snavely, Steven M Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. In ACM siggraph 2006 papers, pages 835–846. 2006. 2
- [71] Noah Snavely, Steven M Seitz, and Richard Szeliski. Modeling the world from internet photo collections. *International journal of computer vision*, 80:189–210, 2008. 2
- [72] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of rgb-d slam systems. In 2012 IEEE/RSJ international conference on intelligent robots and systems, pages 573–580. IEEE, 2012. 6
- [73] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. 2
- [74] Jiaming Sun, Yiming Xie, Linghao Chen, Xiaowei Zhou, and Hujun Bao. Neuralrecon: Real-time coherent 3d reconstruction from monocular video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15598–15607, 2021. 2, 3
- [75] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [76] Chengzhou Tang and Ping Tan. Ba-net: Dense bundle adjustment network. arXiv preprint arXiv:1806.04807, 2018.
 2

- [77] Keisuke Tateno, Federico Tombari, Iro Laina, and Nassir Navab. Cnn-slam: Real-time dense monocular slam with learned depth prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6243–6252, 2017. 2
- [78] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. Advances in neural information processing systems, 34:16558–16569, 2021. 2, 6
- [79] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. Advances in Neural Information Processing Systems, 36, 2024. 2
- [80] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment—a modern synthesis. In Vision Algorithms: Theory and Practice: International Workshop on Vision Algorithms Corfu, Greece, September 21–22, 1999 Proceedings, pages 298– 372. Springer, 2000. 2
- [81] Hengyi Wang and Lourdes Agapito. 3d reconstruction with spatial memory. *arXiv preprint arXiv:2408.16061*, 2024. 3, 5, 6, 7
- [82] Hengyi Wang, Jingwen Wang, and Lourdes Agapito. Coslam: Joint coordinate and sparse parametric encodings for neural real-time slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13293–13302, 2023. 6
- [83] Jianyuan Wang, Nikita Karaev, Christian Rupprecht, and David Novotny. Vggsfm: Visual geometry grounded deep structure from motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21686–21697, 2024. 2
- [84] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689, 2021. 2
- [85] Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang. Moge: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision, 2024. 2
- [86] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 20697–20709, 2024. 2, 3, 5, 6, 7, 8
- [87] 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. 2020. 5
- [88] Philippe Weinzaepfel, Vincent Leroy, Thomas Lucas, Romain Brégier, Yohann Cabon, Vaibhav Arora, Leonid Antsfeld, Boris Chidlovskii, Gabriela Csurka, and Jérôme Revaud. Croco: Self-supervised pre-training for 3d vision tasks by cross-view completion. *Advances in Neural Information Processing Systems*, 35:3502–3516, 2022. 3
- [89] Magnus Wrenninge and Jonas Unger. Synscapes: A photorealistic synthetic dataset for street scene parsing. arXiv preprint arXiv:1810.08705, 2018. 5

- [90] Rundi Wu, Ruiqi Gao, Ben Poole, Alex Trevithick, Changxi Zheng, Jonathan T Barron, and Aleksander Holynski. Cat4d: Create anything in 4d with multi-view video diffusion models. arXiv preprint arXiv:2411.18613, 2024. 4
- [91] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P Srinivasan, Dor Verbin, Jonathan T Barron, Ben Poole, et al. Reconfusion: 3d reconstruction with diffusion priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21551–21561, 2024. 3
- [92] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10371–10381, 2024. 2
- [93] Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1281–1292, 2020. 2
- [94] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer* vision (ECCV), pages 767–783, 2018. 2
- [95] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22, 2023. 5
- [96] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 4578–4587, 2021. 3
- [97] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19447–19456, 2024. 2
- [98] Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the presence of motion. arXiv preprint arxiv:2410.03825, 2024. 3, 5, 6, 7
- [99] 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. 4
- [100] Xiaoshuai Zhang, Sai Bi, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. Nerfusion: Fusing radiance fields for largescale scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5449–5458, 2022. 2, 3
- [101] 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*, pages 20–37. Springer, 2022. 3, 6
- [102] Wang Zhao, Shaohui Liu, Hengkai Guo, Wenping Wang, and Yong-Jin Liu. Particlesfm: Exploiting dense point trajectories for localizing moving cameras in the wild. In *European*

Conference on Computer Vision, pages 523–542. Springer, 2022. 6

- [103] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. ACM Trans. Graph. (Proc. SIGGRAPH), 37, 2018. 5
- [104] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 12786–12796, 2022. 6
- [105] Zihan Zhu, Songyou Peng, Viktor Larsson, Zhaopeng Cui, Martin R Oswald, Andreas Geiger, and Marc Pollefeys. Nicer-slam: Neural implicit scene encoding for rgb slam. In 2024 International Conference on 3D Vision (3DV), pages 42–52. IEEE, 2024. 2