

MVSAnywhere: Zero-Shot Multi-View Stereo

Sergio Izquierdo^{3,*} Mohamed Sayed¹ Michael Firman¹ Guillermo Garcia-Hernando¹
 Daniyar Turmukhambetov¹ Javier Civera³ Oisin Mac Aodha² Gabriel Brostow^{1,4} Jamie Watson^{1,4}
¹Niantic ²University of Edinburgh ³I3A, Universidad de Zaragoza ⁴UCL
<https://nianticlabs.github.io/mvsanywhere/>

Abstract

Computing accurate depth from multiple views is a fundamental and longstanding challenge in computer vision. However, most existing approaches do not generalize well across different domains and scene types (e.g. indoor vs. outdoor). Training a general-purpose multi-view stereo model is challenging and raises several questions, e.g. how to best make use of transformer-based architectures, how to incorporate additional metadata when there is a variable number of input views, and how to estimate the range of valid depths which can vary considerably across different scenes and is typically not known a priori? To address these issues, we introduce MVSA, a novel and versatile Multi-View Stereo architecture that aims to work Anywhere by generalizing across diverse domains and depth ranges. MVSA combines monocular and multi-view cues with an adaptive cost volume to deal with scale-related issues. We demonstrate state-of-the-art zero-shot depth estimation on the Robust Multi-View Depth Benchmark, surpassing existing multi-view stereo and monocular baselines.

1. Introduction

Estimating accurate depth from multiple RGB images is a core challenge in 3D vision, and a building block for downstream applications like 3D reconstruction and autonomous driving. Recent approaches in learning-based multi-view stereo (MVS) are capable of generating accurate depths [6, 89, 90]. However, existing methods typically struggle to generalize to scene and camera setups that differ significantly from those in their training data. As a result, there is a pressing need for general-purpose MVS methods that are more robust to differences between the training and test distributions.

We take inspiration from the recent explosion in scene-agnostic *single-view* depth models, which predict plausible metric [2, 25, 29, 54, 76, 92] or up-to-scale [38, 56, 87, 88] depth using only a single image as input. These models are

*Work done during an internship at Niantic.




Figure 1. Our **MVSA** model results in high-quality reconstructions from posed images, and is superior to existing monocular and MVS methods. Here we compare with Depth Pro [3], a recent monocular method which produces sharp and good looking depth maps, but can have inconsistent scaling of depths, which are required for good meshes. We also include a variant of MAST3R [42] that we have augmented with ground truth camera poses. Our model gives sharp depth maps which are also accurate and 3D consistent, producing high-quality meshes in zero-shot environments.

typically trained on large curated sets of synthetic and/or real RGB-D data, endowing them with impressive generalization performance on previously unseen data. Single-view models are, however, inherently limited by their input.

For our specific depth prediction target, constraining the model’s input to just one image forces it to use single-view geometry cues (*e.g.* vanishing points) and learned patterns [14], while losing the stronger multi-view signal. While there are temporal extensions of these single view models [53, 64, 75, 80, 85], their focus is on temporal perceptual consistency, and not necessarily multi-view consistency. In application contexts where multiple views are available at inference time, it stands to reason that these lead to significantly more accurate depth estimates [35, 47, 66].

Developing a general-purpose MVS method, however, raises two significant challenges. Firstly, it should be able to deal with arbitrary depth ranges. Existing MVS methods typically require a known range of depths to ‘search’ over along epipolar lines, corresponding to a discrete set of depth bins used to build a cost volume. These depths are typically either fixed (and chosen from the range of depths in the training data) [59] or are provided at test time for each image [71, 89, 90]. Secondly, the many emerging benefits of ViTs [16] motivate us to find a way to ‘upgrade’ parts of standard MVS architectures that are still CNNs.

To address these challenges, we introduce a new general-purpose MVS method named Multi-View Stereo Anywhere (MVSA). Similarly to recent performant monocular methods, it is trained on a large and diverse set of data, spanning diverse depth ranges. Along with harmonizing these training signals, our main technical contributions are:

- A novel transformer-based architecture that processes the multi-view cost volume, while *also* incorporating monocular features. We propose a Cost Volume Patchifier that tokenizes the cost volume without losing its details, while also incorporating features from a monocular ViT.
- We propose a view-count-agnostic and scale-agnostic mechanism to construct the cost volume using geometric metadata given any number of input source frames. This is in contrast to the established practice [59] of concatenating geometric metadata from a fixed number of frames to build the cost volume.

MVSA predicts highly accurate and 3D-consistent *depths*, obtaining state-of-the-art results on the Robust Multi-View Depth Benchmark [63], which contains a variety of challenging held-out datasets. We also report scores for some new single- and multi-view methods for comparison. Our better depths result in improved 3D mesh *reconstruction* compared to alternative depth-based reconstruction methods (Fig. 1). Code and pretrained models are available at <https://github.com/nianticlabs/mvsanywhere>.

2. Related Work

Multi-view stereo (MVS). MVS algorithms estimate depth from posed multi-view images using epipolar geometry [72]. Given calibrated cameras, early methods estimated depth by matching image patches [21, 61]. Subsequently,




Figure 2. **MVS datasets cover a wide range of depth values.** Here we show the distribution of % depths in the DTU [36], ScanNet [12], ETH3D [62], Tanks and Temples [41], and KITTI [23] datasets, as a stacked bar chart. Note the log x-axis. This wide range of depth values can be challenging when it comes to constructing meaningful cost volumes and predicting the final depths.

deep learning approaches were introduced, first for stereo matching [94] and later improved via end-to-end learning, typically using plane-sweep cost volumes [10, 24, 33, 34, 39, 69, 74, 90, 95]. Subsequent methods introduced advances in architectures [6, 7, 15], increased robustness to occlusion and moving objects [46, 83, 97], integrated temporal information [17], improved model efficiency [59, 93], jointly estimated camera pose [42, 77] and ingested prior geometry estimates to improve depths [60].

With some exceptions [96], earlier stereo and MVS methods were traditionally both trained and tested on the same dataset/domain, and were limited in their ability to generalize to out-of-distribution data. This domain generalization issue is a consequence of most performant learning-based MVS methods being data-hungry. Approaches such as a training on synthetic [48, 98] or pseudo-labeled depth [30] can be effective, but so far, struggle to span a diverse range of scene types and scales. Self-supervised approaches can be trained without depth supervision, but current methods produce inferior depths compared to fully supervised approaches [13, 40, 86]. Concurrent with our work, [26, 81] trained large *binocular* stereo models on large synthetic datasets.

Adaptive cost volumes. One of the challenges in developing a general-purpose domain-agnostic MVS method is that different scenes can contain wildly different depth ranges, *e.g.* indoor scenes are limited to a few meters, while outdoor ones can span much larger distances. This is a problem as conventional cost volumes require a known depth range, which is typically just estimated based on the minimum and maximum depth values in the training set. As a result, there is a need for cost volumes that are not restricted to a pre-defined range or bins, and instead are adaptive. In the context of self-supervised learning with unscaled poses, [80] estimated bin ranges at training time via an exponential moving average of the depth predictions. Another approach is to predict bin centers iteratively in a coarse-to-fine manner, where the outputs from the previous iteration are used

to seed the range in the next [24, 49, 99]. Alternatively, the bin offsets can be predicted by a learned network [11] or from estimated depth uncertainty [9, 43]. We estimate cost volume depth ranges to enable us adapt to any range of depths, while prior work has done this when the test time range is known, but they wish to reduce computation or enhance detail.

Single-view depth. Monocular depth methods, trained using supervised learning, do not have access to multi-view images at inference time so instead aim to estimate a depth map from a *single* image [18–20, 45]. Building on highly advanced and effective image backbones [52], more recent monocular methods have focused on making *general-purpose* depth estimation models, which aim to work on arbitrary scenes [8, 56]. Further works have scaled up the size of models and datasets, training on combinations of real and/or synthetic data [87, 88], and have used stronger image-level priors [27, 38]. One of the limitations of models trained from stereo-image-derived supervision without known baselines [56] or human annotations [8] is that these only enable a relative, and not metric (*e.g.* in meters), depth prediction.

Other monocular models predict *metric* depth [2, 3, 29, 76, 92]. Not only does this rely on appropriate training data, but also requires an understanding of camera intrinsics, which are often a required additional input to the network. Conventional monocular methods are inherently limited by only incorporating information from single views at inference time, even when multi-view information is available [80]. On the other hand, with recent advances, they can still provide a very valuable signal when only one image is available. As in [1, 79, 84], we combine features extracted from a monocular depth model with a multi-view cost volume to better leverage monocular and multi-view cues.

3. General-purpose Multi-View Stereo

Our model takes as input a $H \times W$ reference image I_r together with neighboring source frames $I_{i \in \{1 \dots N\}}$, each with their relative poses and intrinsics. At test time we aim to predict a dense depth map \hat{D}_r for I_r . For ours to be a general-purpose MVS method, we seek to:

1. **Generalize to any domain.** Most current MVS methods are typically trained on and tested on data from similar domains, *e.g.* indoor only or driving only.
2. **Generalize to any range of depths.** Predicted depth maps need to be accurate for nearby surfaces (*e.g.* for robotics) or for more distant ones (*e.g.* for drones and autonomous driving). In some scenarios like SfM, the depths and camera poses are in a non-metric up-to-scale coordinate system. Hence, general-purpose MVS should be robust to the scale of the coordinate system.
3. **Be robust to the number and selection of source frames.** Traditional MVS systems can struggle when




Figure 3. **Our general-purpose multi-view depth estimation model.** We start with a cost-volume based architecture, which matches deep features between views at different hypothesized depths. Key for performance are our Cost Volume Patchifier and Mono/Multi Cue Combiner. These also fuse single-view information coming from the Reference Image Encoder and source views.

there is little overlap between source and reference frames. We also want MVS methods to be agnostic to the number of source frames available at test time.

4. **Predict 3D-consistent depths.** Depths from one viewpoint should be consistent with those predicted from different viewpoints. Fusion of consistent depth maps will produce a mesh with accurate estimates of 3D surfaces.

While prior works have tackled these problems in turn, we are the first model, to the best of our knowledge, to tackle all four problems in a single system.

3.1. MVSAnywhere

We introduce **MVSAnywhere** (MVSA), a novel general-purpose MVS system which is designed to embody each of the previous properties. To help us learn from diverse datasets and hence **generalize to any domain**, we use a large transformer-based architecture, which takes as input: (1) multi-view information from the reference and source images, and (2) single-view information, which is extracted directly from the reference image via a monocular *reference image encoder*. The overall architecture (Fig. 3 and supplementary) is broadly inspired by recent MVS approaches, *e.g.* [59]. It comprises five key components:

Feature extractor. This encodes the source and reference images into deep feature maps \mathcal{F}_r and $\mathcal{F}_{i \in \{1 \dots N\}}$, that will be processed via a cost volume. We use the first two blocks of a ResNet18 [28] for this encoder, producing feature maps at resolution $H/4 \times W/4$.

Cost volume. Following *e.g.* [7, 33, 39, 74], we warp feature maps \mathcal{F}_i from each source view to the reference one using a set of hypothesized depth values (*i.e.* bins) \mathcal{D} . We then concatenate these warped features and \mathcal{F}_r with appropriate *metadata*, following [59]. See Sec. 3.2 for our specific novel contributions in this matter.

Reference image encoder. This extracts powerful deep monocular features for I_r . We use the ViT Base [16] encoder from Depth Anything V2 [88], with their pretrained weights for relative monocular depth estimation, which help us to be **robust to limited overlaps** between source

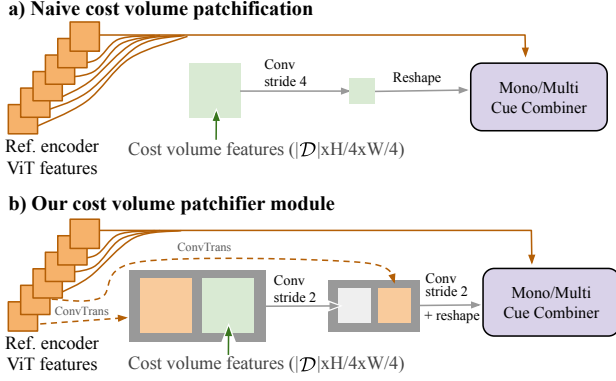


Figure 4. **Our cost volume patchifier** enables high-quality information to be extracted from a $|\mathcal{D}| \times \frac{H}{4} \times \frac{W}{4}$ cost volume, ready for input to the Mono/Multi Cue Combiner ViT. (a) Shows the naive approach to patchification. (b) Our approach makes better use of the reference image features.

and reference frames. As ViT Base operates on 14×14 patches, the reference image is resized to $\frac{14H}{16} \times \frac{14W}{16}$ resolution before feeding to ViT Base, such that the extracted features are size $\frac{H}{16} \times \frac{W}{16}$.

Mono/Multi Cue Combiner. This converts the “patchified” features of the cost volume and reference image into a sequence of features which go to our depth decoder. Monocular and multi-view cues are combined by a novel component described in Sec. 3.3.

Depth Decoder. Based on the decoder from [55], MVSA progressively upsamples and processes features from the Mono/Multi Cue Combiner module to produce the final depth map at the reference image resolution.

3.2. Metadata Agnostic to View Count and Scale

SimpleRecon [59] demonstrated that readily available metadata, *e.g.* geometric and camera pose information, can be incorporated into the cost volume to improve depths. For each pixel location (u_r, v_r) in I_r and depth bin k in \mathcal{D} , we back-project the pixel to a 3D point P and then reproject it into every source view I_i . The specific metadata for the bin with coordinates (u_r, v_r, k) in the cost volume includes: the dot product between feature vector $\mathcal{F}_r(u_r, v_r)$ and corresponding pixels in $\mathcal{F}_i(u_i, v_i)$, ray directions from source and reference origin to P , depths in reference and source views, angle between rays from reference and source views, relative poses between the reference and source cameras, and depth validity masks (in case P is outside a source frame frustum). See our supplementary for full details.

SimpleRecon [59]’s cost volume concatenates metadata from all eight source frames and runs an MLP to produce one single cost (matching score) per spatial location and depth hypothesis. While this gives good scores, its limitation is that it requires *exactly* eight source frames

for every training and test reference image, limiting the model’s flexibility (note though that traditional MVS methods are typically already view-count agnostic). To address this limitation, we introduce a *view-count-agnostic metadata* component which enables a single model to **generalize to any number of source views**. For each source frame, we run an MLP that ingests the metadata from the reference frame and the source frame and predicts two values: a score and a weight. This results in N scores and N weights. A weighted sum of the N scores is computed after the N weights go through softmax. This weighted sum is used as the value in the cost volume at every pixel location (u, v) and depth hypothesis k . Our novel module enables aggregation of the matching score and confidence for each source frame, while allowing for a variable number of source frames for each I_r .

The source camera poses may be close to the reference, or far from it. To be more invariant to this possible range of scales, we also make the metadata **scene scale-agnostic**. To this end, we normalize the relative pose measures of the metadata using a maximum across all the source frames for a given reference frame. We also normalize the depth hypothesis metadata using the maximum and minimum of \mathcal{D} .

As the scene scale information is not provided to the rest of our network, we rescale the depth predictions to match the scale of the input poses. Our depths are predicted with a sigmoid function σ over the logit x . To align the prediction of the network with the cost volume, the sigmoid output is scaled by the depth range of the cost volume, so

$$\hat{D}_r = \exp(\log(d_{\min}) + \log(d_{\max}/d_{\min}) \cdot \sigma(x)). \quad (1)$$

3.3. Mono/Multi Cue Combiner

Given the cost volume of shape $|\mathcal{D}| \times \frac{H}{4} \times \frac{W}{4}$, and the reference image encoder features of shape $C \times \frac{H}{16} \times \frac{W}{16}$ (outputs of different blocks of the reference image encoder), we pose the question: how can we best combine these features to provide the strongest signal for the decoder? Motivated by the recent success of transformer architectures in single-view depth prediction [3, 55, 88], we use a ViT-Base network to process these features in a *Mono/Multi Cue Combiner* network, which produces a sequence of tokens for the decoder to transform into a depth prediction.

To effectively achieve this we need to i) convert the cost volume into a token sequence without sacrificing information and ii) incorporate the monocular cues to help in decoding sharp depth. For i), a naive approach would be to apply a strided convolution projecting to the ViT token dimensions, resembling how RGB images are patchified. However this is suboptimal, for it lacks contextual information on how to achieve this downsampling. Instead, we propose a **cost volume patchifier** module. This guides the downsampling process with information from the first two blocks of the

Name	Scenes	#total scenes	# total images	# training tuples	Metric poses?	Moving objects?
Hypersim [58]	Indoor	461	77K	45K	Yes	No
TartanAIR [78]	Indoor, Outdoor	30	1M	92K	Yes	Yes
BlendedMVG [91]	Indoor, Outdoor, Aerial	389	110K	97K	No	No
MatrixCity [44]	Outdoor, Aerial	1	519K	40K	Yes	No
VKITTI2 [5, 22]	Outdoor	5	21K	40K	Yes	Yes
Dynamic Replica [37]	Indoor	484	145K	70K	Yes	Yes
MVSSynth [33]	Outdoor	117	12K	3K	No	Yes
SAIL-VOS 3D [31, 32]	Indoor, Outdoor	6807	237K	21K	Yes	Yes

Table 1. We train on eight MVS datasets from a variety of domains. All these datasets are synthetically rendered, giving them perfect ground truth depths and camera calibration. However, BlendedMVG uses real textures on their assets.

reference image encoder. We convert the cost volume into tokens using two strided convolutions, but first, concatenate each of them with the monocular features of the first two blocks, transposed and projected at 1/4 and 1/8 of the input resolution, respectively. The output of this module is a sequence of $\frac{H}{16} \times \frac{W}{16}$ tokens, matching the sequence length of the monocular features. These tokens are then fed into a ViT-B initialized with DINOv2 weights (see Fig. 4).

For ii) we add the tokens from the cost volume with the ones from the reference image encoder after projecting the latter with a linear layer. We repeat this process at blocks 2, 5, 9, and 11 of the ViT to incorporate multiple levels of monocular cues. This simple mechanism allows the network to refine and regularize the cost volume with the help of the reference image structure.

3.4. Generalizing to Any Range of Depths

When building a cost volume, a set of depth hypotheses (*i.e.* bins) \mathcal{D} are used to warp feature maps \mathcal{F}_i to I_r . This raises the question: How do we choose \mathcal{D} to **generalize to any range of depths**? Depth ranges vary hugely across datasets (see Fig. 2), so using the same fixed range is suboptimal.

We address this with a cascaded cost volume approach, first introduced in 3D stereo matching [24, 49, 99]. While these works start from a known ‘ground truth’ depth range, we use the known intrinsics and extrinsics to infer the minimum and maximum depths that could be matched between I_r and each I_i . We space our initial depth bins logarithmically within this range, then make an initial depth prediction. The min and max values of this initial estimate are then used to rebuild the cost volume for a final depth prediction. This iterative process occurs only at test time; during training, we use the known depth range. Full details are in the supplementary. Importantly, previous methods that are provided with an exact depth range learn to predict depths that cover all the depth hypotheses. Thus, when using a rough estimate of the range, these methods fail to align the prediction to the actual valid depths. To further mitigate this issue, we augment the ground truth ranges via a random perturbation during training.

3.5. Implementation Details

Losses. We use the supervised losses from [59]. These comprise an L1 loss between the log of the ground truth and the log of the predicted depth values, and a gradient and normals loss. Training losses are applied to four output scales of the decoder. At inference, only the final largest-scale prediction is used. We take as input 640×480 images, and output depth maps at the same resolution. We use 64 depth bins in \mathcal{D} sampled in log space.

Keyframes. For datasets with dense sequences, we choose reference and source frames with the strategy of [17, 59]. To be robust to sparser sets of frames, we also select tuples based on geometry overlap, obtaining tuples of not necessarily consecutive frames. Full details on our architecture and training strategy are provided in the supplementary.

Training data. For MVSA to **generalize across domains**, we train on a large and diverse set of synthetic datasets, as listed in Tab. 1. A subset of these training datasets contain moving objects. Our reference image encoder is initialized from Depth Anything V2 (DAV2) [88], which uses a teacher network trained on synthetic datasets similar to ours, and a student network distilled using various real images that do not overlap with our evaluation benchmarks. DAV2 was initialized from a pretrained DINOv2 [52] network, in turn trained on internet images.

4. Experiments

We evaluate MVSA on both depth estimation and 3D reconstruction tasks. We also implement and report scores for a set of new baselines, which have never before been evaluated on the benchmarks we use.

4.1. Baselines

Where possible, we obtain results directly from prior works [63, 77]. We also evaluate and implement other strong baselines that did not previously report performance on diverse MVS benchmarks. These include: (i) A strong monocular baseline in the form of DAV2 [88]. To account for the unknown affine transform, we align its predictions to the ground truth using least squares. (ii) MAST3R [42] (raw depth estimate) which involves passing the reference and one other source image as input and taking the z component of the point cloud as the depth prediction. (iii) MAST3R (plus our triangulation) which is a novel extension of MAST3R so that it can use provided extrinsics and intrinsics, when available. For each of the available source images, we use MAST3R descriptors to match points with the reference image. We then triangulate points from such matches, rescale the raw depth predictions, and aggregate the point clouds from the different views using a sum weighted by the predicted confidences. Note, this method requires one forward pass and thousands of triangulations

Approach	GT	GT	Align	KITTI		ScanNet		ETH3D		DTU		T&T		Average		
	Poses	Range		rel ↓	τ ↑	rel ↓	τ ↑	rel ↓	τ ↑	rel ↓	τ ↑	rel ↓	τ ↑	rel ↓	τ ↑	time [s] ↓
a) Depth from frames (w/o poses)																
DeMoN [70]	✗	✗	t	15.5	15.2	12.0	21.0	17.4	15.4	21.8	16.6	13.0	23.2	16.0	18.3	0.08
DeepV2D KITTI [69]	✗	✗	med	(3.1)	(74.9)	23.7	11.1	27.1	10.1	24.8	8.1	34.1	9.1	22.6	22.7	2.07
DeepV2D ScanNet [69]	✗	✗	med	10.0	36.2	(4.4)	(54.8)	11.8	29.3	7.7	33.0	8.9	46.4	8.6	39.9	3.57
MAST3R [42] (raw output)	✗	✗	med	3.3	67.7	(4.3)	(64.0)	2.7	79.0	3.5	66.7	(2.4)	(81.6)	3.3	71.8	0.07
MAST3R [42] (raw output)	✗	✗	✗	61.4	0.4	(12.8)	(19.4)	43.8	3.1	145.8	0.5	(66.9)	(0.0)	66.1	4.7	0.07
b) Depth from frames and poses (with per-image range provided)																
MVSNet [90]	✓	✓	✗	22.7	36.1	24.6	20.4	35.4	31.4	(1.8)	(86.0)	8.3	73.0	18.6	49.4	0.07
MVSNet Inv. Depth [90]	✓	✓	✗	18.6	30.7	22.7	20.9	21.6	35.6	(1.8)	(86.7)	6.5	74.6	14.2	49.7	0.32
Vis-MVSNet [97]	✓	✓	✗	9.5	55.4	8.9	33.5	10.8	43.3	(1.8)	(87.4)	4.1	87.2	7.0	61.4	0.70
PatchmatchNet [71]	✓	✓	✗	10.8	45.8	8.5	35.3	19.1	34.8	(2.1)	(82.8)	4.8	82.9	9.1	56.3	0.28
MVSFormer++ DTU+BlendedMVG [7]	✓	✓	✗	4.4	65.7	7.9	39.4	7.8	50.4	(0.9)	(95.3)	3.2	88.1	4.8	67.8	0.78
MVSFormer++ CE our data [7]	✓	✓	✗	4.4	63.9	6.4	43.3	6.7	56.4	(1.2)	(90.7)	2.6	88.3	4.3	68.5	0.78
MVSFormer++ regression our data [7]	✓	✓	✗	3.7	67.2	5.7	45.1	5.3	57.5	(1.2)	(90.5)	2.2	88.6	3.6	69.8	0.78
c) Single-view depth																
Depth Pro [3] †	✗	✗	med	6.1	39.6	(4.3)	(58.4)	6.1	53.5	5.6	49.6	5.6	57.5	5.6	51.7	5.16
Depth Pro [3] †	✗	✗	✗	13.6	14.3	9.2	19.7	28.5	8.7	161.8	3.5	38.3	4.4	50.3	10.1	5.16
Metric3D [29] †	✗	✗	med	5.1	44.1	2.4	78.3	4.4	54.5	10.1	39.5	6.2	48.0	5.6	52.9	0.46
Metric3D [29] †	✗	✗	✗	8.7	13.2	6.2	19.3	12.7	13.0	890.5	1.4	16.7	13.7	187.0	12.1	0.46
UniDepthV2 [54] †	✗	✗	med	4.0	55.3	(2.1)	(82.6)	3.7	66.2	3.2	72.3	3.6	68.4	3.3	68.9	0.29
UniDepthV2 [54] †	✗	✗	✗	13.7	4.8	(3.2)	(61.3)	15.4	11.9	964.8	1.3	16.7	12.7	202.7	18.4	0.29
UniDepthV1 [54] †	✗	✗	med	4.4	51.6	(1.9)	(84.3)	5.4	48.4	9.3	31.8	9.6	38.7	6.1	51.0	0.21
UniDepthV1 [54] †	✗	✗	✗	5.2	39.5	(2.7)	(69.4)	48.2	1.8	583.3	1.0	30.7	4.2	134.0	23.2	0.20
DepthAnything V2 (ViT-B) [88]	✗	✗	lstsq †	6.6	38.6	4.0	58.6	4.7	56.5	2.6	74.7	4.5	57.5	4.8	54.1	0.05
d) Depth from frames and poses (w/o per-image range)																
Fast-MVSNet [93]	✓	✗	✗	12.1	37.4	287.1	9.4	131.2	9.6	(540.4)	(1.9)	33.9	47.2	200.9	21.1	0.35
MVS2D ScanNet [89]	✓	✗	✗	73.4	0.0	(4.5)	(54.1)	30.7	14.4	5.0	57.9	56.4	11.1	34.0	27.5	0.05
MVS2D DTU [89]	✓	✗	✗	93.3	0.0	51.5	1.6	78.0	0.0	(1.6)	(92.3)	87.5	0.0	62.4	18.8	0.06
Robust MVD Baseline [63]	✓	✗	✗	7.1	41.9	7.4	38.4	9.0	42.6	2.7	82.0	5.0	75.1	6.3	56.0	0.06
MAST3R (plus our triangulation)	✓	✗	✗	3.4	66.6	(4.5)	(63.0)	3.1	72.9	3.4	67.3	(2.4)	(83.3)	3.4	70.1	0.72
MVSA (Ours)	✓	✗	✗	3.2	68.8	3.7	62.9	3.2	68.0	1.3	95.0	2.1	90.5	2.7	77.0	0.12

Table 2. We set a new SOTA in depth estimation on the RMVDB. See Sec. 4 for details of the metrics, baselines and groupings. A full version of this table, including older baselines, appears in the supplementary. Monocular methods with † are given ground truth intrinsics. The best result for each section appears in **bold**, and (parentheses) indicate results where the evaluation dataset is in the training set.

per source view, significantly reducing its speed. MAST3R trains on ScanNet [12] and MegaDepth [45] (which contains a subset of the Tanks and Temples dataset [41]). Baseline implementation details are in the supplementary.

Benchmark. We evaluate ‘zero-shot’ depth estimation performance on the five multi-view datasets from the RMVDB benchmark [63], which are not included in our training data. It contains the KITTI [23] ScanNet [12], ETH3D [62], DTU [36], and Tanks and Temples [41] datasets and represents a diverse set of evaluation scenarios, *e.g.* driving sequences, room scans, building scans, and tabletop objects, among others. We use the evaluation procedure and source view selection procedure from [63], allowing direct comparison to previous approaches.

Methods are grouped into four different types (a-d) depending on the information they are provided, *e.g.* if they are given GT cameras, GT depth ranges, *etc.* MVSA naturally fits into type (d), where all methods are given GT poses, so need to predict depth directly in metric scale and hence do not need any alignment or knowledge of the GT depth range. Note, some methods train on the training splits

of one, or more, of the benchmark datasets, thus achieving very high scores in those cases. We denote these in Tab. 2 with a parenthesis around them.

Metrics. We report two commonly used metrics to compare the predicted \hat{d} and GT depth d . The absolute relative depth (rel) is computed per-pixel as $|\hat{d} - d|/d$, while the inlier percentage τ , with threshold 1.03, is computed per-pixel as $[\max(\hat{d}/d, d/\hat{d}) < 1.03]$, where $[\]$ is the Iverson bracket. Both metrics are averaged over all valid GT pixels in each test image, before averaging over all images.

Results. Tab. 2 depicts the quantitative results, where we outperform all baselines across most metrics. Qualitative results in Fig. 7 demonstrate that our MVSA model produces depth maps with superior edge detail and consistent scaling across a variety of scenes, visually outperforming prior methods. MVSA also performs well on moving objects, *e.g.* as found in KITTI; see also Fig. 6. Both **MAST3R triangulated** and the **Robust MVD Baseline** exhibit poor edge quality, limiting their suitability for applications such as single-image novel view synthesis [65], which requires sharp depth boundaries. While Depth Pro produces

Approach	ScanNet		ETH3D	
	rel ↓	τ ↑	rel ↓	τ ↑
Robust MVD Baseline [63]	6.02	47.83	5.75	71.64
MAST3R Triangulated	(3.88)	(68.68)	2.37	84.90
MVSA (Ours)	3.22	69.45	1.27	93.24

Table 3. **Our variant of RMVDB.** We use better test-time tuples for ScanNet, and for ETH3D we use the undistorted test images.




Figure 5. **Many MVS models fail in areas of poor frame overlap.** Here we show how MVSAFormer++ (right) fails to recover geometry in areas of the image where there are no matching pixels between source and target views (see the top left corner). Our model (middle) handles this situation gracefully.

sharp edges, it frequently displays incorrect depth scaling. In contrast, our MVSA model combines competitive quantitative performance with sharper edges, making it ideal for tasks that demand both visual and depth accuracy. Finally, GT depth-based median and least squares scaling of monocular methods and depth from frames methods (w/o poses) is crucial for good scores, while MVSA consistently predicts high-quality and metric depths.

Alternative Variant of RMVDB. We further evaluate some of the leading models on a RMVDB variant, in which we change some conditions to better reflect real-world scenarios. In this variant, for ScanNet we use keyframes using the strategy of [17], rather than the temporally sequential keyframes provided by the benchmark. For ETH3D we undistort both the images and the ground truth using their provided Thin-Prism [82] camera parameters. Results are shown in Tab. 3. On this revised benchmark, we more comprehensively outperform the baselines.

4.2. Ablations Study

In Tab. 4 we validate our design decisions by turning on and off sections of our system. We train all ablations at a smaller resolution (512×384 input), and without using metadata, for efficiency. At this resolution, Row **A** is ‘ours’ and all other rows are ablations relative to this. Row **B** replaces our standard ViT-B with the smaller ViT Small, both for the cost volume ViT and the reference image encoder. The reference image encoder is initialized from Depth Anything v2 (small). Row **C** uses our training data and pipeline, but with the fully-convolutional architecture from SimpleRecon [59] (without metadata). Row **D** is our system but without adding noise to the ground truth range at training time (Section 3.4). Although this method can excel when the initial range is accurate, it can fail to generalize (see DTU). Row **E** is our full architecture but without the pretrained en-




Figure 6. We handle **dynamic objects** significantly better than traditional MVS *e.g.* MVSAFormer++. See supplementary for details of the MVSAFormer++ variants shown here.

	KITTI		ScanNet		ETH3D		DTU		T&T		Average	
	rel ↓	τ ↑	rel ↓	τ ↑	rel ↓	τ ↑	rel ↓	τ ↑	rel ↓	τ ↑	rel ↓	τ ↑
A Ours (no metadata, low res.)	3.39	66.88	3.86	60.82	3.11	70.17	2.43	92.05	2.23	88.38	3.00	75.66
B w/ ViT Small	3.57	64.34	4.40	56.57	3.63	64.61	2.69	91.52	2.71	84.51	3.40	72.31
C w/ [59]’s architecture	3.63	65.94	5.03	51.76	3.74	63.09	1.77	90.97	2.78	87.90	3.39	71.93
D w/o noise on GT range	3.33	66.83	5.14	52.77	3.53	66.32	13.45	89.21	2.34	87.64	5.32	74.89
E w/o DAV2 weights	3.45	65.42	4.58	57.14	3.48	65.59	2.14	92.48	2.56	86.01	3.24	73.33
F w/ fixed bins [0-100m]	3.41	64.33	3.80	61.62	3.15	67.20	4.11	65.64	2.36	85.72	3.37	68.90
G no MMCC ViT	3.54	65.32	4.39	56.94	3.56	65.27	3.07	90.49	2.46	87.54	3.41	73.11
H w/o bin refinement	3.57	63.39	5.18	51.05	3.50	67.93	6.80	82.12	2.27	87.04	4.26	70.31
I Naive patchify	3.66	62.61	4.27	58.86	3.18	67.27	1.95	91.69	2.46	86.52	3.11	73.39

Table 4. **Ablation Study.** Here we validate our design decisions on RMVDB [63] by ablating various components. See Sec. 4.2 for details. **First** and **second** best scores are indicated.

coder weights from [88]. Instead we initialize with DINOv2 weights. Row **F** is our system without a cascaded cost volume, and instead uses a fixed set of depth bins, losing the ability to refine depth bins and work with arbitrary scales or scene sizes. Row **H** is our full model, but where we take the first depth prediction from the model as our final output, without re-building the cost volume. Even though these bins capture the full range of depths in the test datasets (Fig. 2), we see that performance degrades. Row **G** uses CNN layers instead of ViT to combine mono/multi features. Row **I** uses naive patchification to preprocess the cost volume for input to the mono/multi cue combiner ViT, as outlined in Fig. 4. These results confirm our design decisions, and full implementation details are in the supplementary.

Robustness to pose rescaling. In the supplementary we include results where we scale ScanNet poses $\times 100$, and show our depths are robust to this scaling (*vs.* ours w/o normalization of the metadata, which performs badly).

4.3. Meshing and 3D Reconstruction

To judge the 3D-consistency of our predictions we evaluate our model on ScanNet Mesh Evaluation benchmark using the protocol defined in [4] which also uses source frame selection from [17]. The benchmark uses a GT mesh collected with an active RGBD sensor captured in a video. The evaluation computes point-to-point vertex error from GT to predicted (as accuracy), from predicted to GT (as completion), and the average of the two (as chamfer). Additionally, 200k points are sampled uniformly over each mesh and point-to-point errors thresholded at 5cm distance are used to compute precision, recall and F-score. Almost all competing methods are trained on ScanNet, however our method that was not trained on ScanNet performs comparatively, outperforming many of the methods in Table 5. See Fig. 1 and the supplementary for qualitative meshing results.




Figure 7. **Qualitative comparison of depth prediction results across multiple datasets** (KITTI, ScanNet, ETH3D, DTU, and Tanks & Temples). Rows show different methods: Depth Pro [63], rMVD baseline [63], MAST3R (Triangulated) [42], and our MVSA model, along with RGB inputs (I_r) and ground-truth depths (GT). Depth Pro provides sharp edges but often misestimates depth scale, while our MVSA model captures finer details than MAST3R and rMVD. Depth maps are normalized to ground truth depth range for consistent visualization; see the supplementary material for unnormalized results, especially highlighting scale discrepancies in Depth Pro.

		Comp↓	Acc↓	Chamfer↓	Prec↑	Recall↑	F-Score↑
DeepVideoMVS [17]	S	10.68	6.90	8.79	0.541	0.592	0.563
ATLAS [51]	S	7.16	7.61	7.38	0.675	0.605	0.636
NeuralRecon [68]	S	5.09	9.13	7.11	0.630	0.612	0.619
3DVNet [57]	S	7.72	6.73	7.22	0.655	0.596	0.621
TransformerFusion [4]	S	5.52	8.27	6.89	0.728	0.600	0.655
VoRTX [67]	S	4.31	7.23	5.77	0.767	0.651	0.703
SimpleRecon [59]	S	5.53	6.09	5.81	0.686	0.658	0.671
COLMAP [61]		10.22	11.88	11.05	0.509	0.474	0.489
MAST3R [42] (raw depth)	S+	12.35	12.69	12.52	0.265	0.283	0.272
MAST3R [42] (+ triangulation)	S+	5.38	6.78	6.08	0.572	0.655	0.608
SimpleRecon [59] (trained on our data)		8.07	6.67	7.37	0.501	0.597	0.544
MVSA (Ours)		4.93	6.39	5.66	0.616	0.696	0.652

Table 5. **ScanNet Mesh Evaluation** [4]. Scores adapted from [4, 59]. Rows marked with S were trained on ScanNet only, while those marked S+ were trained on ScanNet and other datasets. Our MVSA model, which was not trained on ScanNet, outperforms many models which were, e.g. [51, 57, 68].

Limitations. While we use multi-view information to generate depths, we do not enforce or encourage temporal

consistency. Techniques for this [30, 47, 85] could work with MVSA. Also, like traditional MVS, our method requires known camera intrinsics and poses; recent works suggest this requirement could be relaxed [50, 73].

5. Conclusions

We introduced MVSAAnywhere, a new general-purpose MVS depth estimation approach. We addressed challenges associated with training on diverse MVS datasets, such as how to best leverage ViT-based architectures, how to incorporate geometric metadata, and how to handle variable depth ranges. Through extensive experimentation, we compare to numerous existing and new baselines. Our contributions result in state-of-the-art zero-shot performance on a range of challenging reconstruction and depth estimation test datasets, in some cases even outperforming models trained on the test domains.

Acknowledgements

We are extremely grateful to Saki Shinoda, Jakub Powierza, and Stanimir Vichev for their invaluable compute infrastructure support. S Izquierdo and J Civera are funded by the Spanish projects PID2021-127685NB-I00 and TED2021-131150B-I00, Aragón project T45_23R, and the scholarship FPU20/02342.

References

- [1] Luca Bartolomei, Fabio Tosi, Matteo Poggi, and Stefano Mattoccia. Stereo anywhere: Robust zero-shot deep stereo matching even where either stereo or mono fail. In *CVPR*, 2025. 3
- [2] 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. 1, 3
- [3] 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. In *ICLR*, 2025. 1, 3, 4, 6, 8
- [4] Aljaz Bozic, Pablo Palafox, Justus Thies, Angela Dai, and Matthias Nießner. TransformerFusion: Monocular RGB scene reconstruction using transformers. *NeurIPS*, 2021. 7, 8
- [5] Yohann Cabon, Naila Murray, and Martin Humenberger. Virtual KITTI 2. *arXiv:2001.10773*, 2020. 5
- [6] Chenjie Cao, Xinlin Ren, and Yanwei Fu. MVSFormer: Multi-view stereo by learning robust image features and temperature-based depth. *TMLR*, 2023. 1, 2
- [7] Chenjie Cao, Xinlin Ren, and Yanwei Fu. MVSFormer++: Revealing the devil in transformer’s details for multi-view stereo. In *ICLR*, 2024. 2, 3, 6
- [8] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. In *NeurIPS*, 2016. 3
- [9] Shuo Cheng, Zexiang Xu, Shilin Zhu, Zhuwen Li, Li Erran Li, Ravi Ramamoorthi, and Hao Su. Deep stereo using adaptive thin volume representation with uncertainty awareness. In *CVPR*, 2020. 3
- [10] Xinjing Cheng, Peng Wang, and Ruigang Yang. Learning depth with convolutional spatial propagation network. *PAMI*, 2019. 2
- [11] Andrea Conti, Matteo Poggi, Valerio Cambareri, and Stefano Mattoccia. Range-agnostic multi-view depth estimation with keyframe selection. In *3DV*, 2024. 3
- [12] 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. 2, 6
- [13] Yuchao Dai, Zhidong Zhu, Zhibo Rao, and Bo Li. Mvs2: Deep unsupervised multi-view stereo with multi-view symmetry. In *3DV*, 2019. 2
- [14] Duolikun Danier, Mehmet Aygün, Changjian Li, Hakan Bilen, and Oisin Mac Aodha. DepthCues: Evaluating monocular depth perception in large vision models. In *CVPR*, 2025. 2
- [15] Yikang Ding, Wentao Yuan, Qingtian Zhu, Haotian Zhang, Xiangyue Liu, Yuanjiang Wang, and Xiao Liu. Transmvsnet: Global context-aware multi-view stereo network with transformers. In *CVPR*, 2022. 2
- [16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2, 3
- [17] Arda Duzceker, Silvano Galliani, Christoph Vogel, Pablo Speciale, Mihai Dusmanu, and Marc Pollefeys. DeepVideoMVS: Multi-view stereo on video with recurrent spatio-temporal fusion. In *CVPR*, 2021. 2, 5, 7, 8
- [18] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015. 3
- [19] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NeurIPS*, 2014.
- [20] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, 2018. 3
- [21] Yasutaka Furukawa and Carlos Hernández. Multi-view stereo: A tutorial. *Foundations and Trends in Computer Graphics and Vision*, 2015. 2
- [22] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR*, 2016. 5
- [23] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *CVPR*, 2012. 2, 6
- [24] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *CVPR*, 2020. 2, 3, 5
- [25] Vitor Guizilini, Pavel Tokmakov, Achal Dave, and Rares Ambrus. GRIN: Zero-shot metric depth with pixel-level diffusion. *arXiv:2409.09896*, 2024. 1
- [26] Xianda Guo, Chenming Zhang, Youmin Zhang, Dujun Nie, Ruilin Wang, Wenzhao Zheng, Matteo Poggi, and Long Chen. Stereo anything: Unifying stereo matching with large-scale mixed data. *arXiv:2411.14053*, 2024. 2
- [27] 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. In *ICLR*, 2025. 3
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3
- [29] 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. *PAMI*, 2024. 1, 3, 6
- [30] 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. In *CVPR*, 2025. 2, 8
- [31] Y.-T. Hu, H.-S. Chen, K. Hui, J.-B. Huang, and A. G. Schwing. SAIL-VOS: Semantic Amodal Instance Level Video Object Segmentation – A Synthetic Dataset and Baselines. In *CVPR*, 2019. 5
- [32] Y.-T. Hu, J. Wang, R. A. Yeh, and A. G. Schwing. SAIL-VOS 3D: A Synthetic Dataset and Baselines for Object Detection and 3D Mesh Reconstruction from Video Data. In *CVPR*, 2021. 5
- [33] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. DeepMVS: Learning multi-view stereopsis. In *CVPR*, 2018. 2, 3, 5
- [34] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. DPSNet: End-to-end deep plane sweep stereo. *ICLR*, 2019. 2
- [35] Sergio Izquierdo and Javier Civera. SfM-TTR: Using structure from motion for test-time refinement of single-view depth networks. In *CVPR*, 2023. 2
- [36] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, 2014. 2, 6
- [37] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. DynamicStereo: Consistent dynamic depth from stereo videos. *CVPR*, 2023. 5
- [38] 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. 1, 3
- [39] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *ICCV*, 2017. 2, 3
- [40] Tejas Khot, Shubham Agrawal, Shubham Tulsiani, Christoph Mertz, Simon Lucey, and Martial Hebert. Learning unsupervised multi-view stereopsis via robust photometric consistency. In *CVPR Workshops*, 2019. 2
- [41] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *Transactions on Graphics*, 2017. 2, 6
- [42] Vincent Leroy, Yohann Cabon, and Jerome Revaud. Grounding image matching in 3D with MAST3R. In *ECCV*, 2024. 1, 2, 5, 6, 8
- [43] Jingliang Li, Zhengda Lu, Yiqun Wang, Jun Xiao, and Ying Wang. Nr-mvsnet: Learning multi-view stereo based on normal consistency and depth refinement. *TIP*, 2023. 3
- [44] Yixuan Li, Lihan Jiang, Linning Xu, Yuanbo Xiangli, Zhenzhi Wang, Dahua Lin, and Bo Dai. MatrixCity: A large-scale city dataset for city-scale neural rendering and beyond. In *ICCV*, 2023. 5
- [45] Zhengqi Li and Noah Snavely. MegaDepth: Learning single-view depth prediction from internet photos. In *CVPR*, 2018. 3, 6
- [46] Xiaoxiao Long, Lingjie Liu, Christian Theobalt, and Wenping Wang. Occlusion-aware depth estimation with adaptive normal constraints. In *ECCV*, 2020. 2
- [47] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *Transactions on Graphics*, 2020. 2, 8
- [48] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, 2016. 2
- [49] Zhenxing Mi, Chang Di, and Dan Xu. Generalized binary search network for highly-efficient multi-view stereo. In *CVPR*, 2022. 3, 5
- [50] Riku Murai, Eric Dexheimer, and Andrew J. Davison. MAST3R-SLAM: Real-time dense SLAM with 3D reconstruction priors. In *CVPR*, 2025. 8
- [51] Zak Murez, Tarrence van As, James Bartolozzi, Ayan Sinha, Vijay Badrinarayanan, and Andrew Rabinovich. Atlas: End-to-end 3D scene reconstruction from posed images. In *ECCV*, 2020. 8
- [52] 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. *TMLR*, 2024. 3, 5
- [53] Vaishakh Patil, Wouter Van Gansbeke, Dengxin Dai, and Luc Van Gool. Don't forget the past: Recurrent depth estimation from monocular video. *Robotics and Automation Letters*, 2020. 2
- [54] 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. 1, 6
- [55] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. 4
- [56] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *PAMI*, 2022. 1, 3
- [57] Alexander Rich, Noah Stier, Pradeep Sen, and Tobias Höllerer. 3DVNet: Multi-view depth prediction and volumetric refinement. In *3DV*, 2021. 8
- [58] 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
- [59] Mohamed Sayed, John Gibson, Jamie Watson, Victor Prisacariu, Michael Firman, and Clément Godard. SimpleRecon: 3D reconstruction without 3D convolutions. In *ECCV*, 2022. 2, 3, 4, 5, 7, 8
- [60] Mohamed Sayed, Filippo Aleotti, Jamie Watson, Zawar Qureshi, Guillermo Garcia-Hernando, Gabriel Brostow, Sara Vicente, and Michael Firman. Doubletake: Geometry guided depth estimation. In *ECCV*, 2024. 2
- [61] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, 2016. 2, 8
- [62] Thomas Schöps, Johannes L. Schönberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and An-

- dreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *CVPR*, 2017. [2, 6](#)
- [63] Philipp Schröppel, Jan Bechtold, Artemij Amiranashvili, and Thomas Brox. A benchmark and a baseline for robust multi-view depth estimation. In *3DV*, 2022. [2, 5, 6, 7, 8](#)
- [64] Jiahao Shao, Yuanbo Yang, Hongyu Zhou, Youmin Zhang, Yujun Shen, Matteo Poggi, and Yiyi Liao. Learning temporally consistent video depth from video diffusion priors. In *CVPR*, 2025. [2](#)
- [65] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3D photography using context-aware layered depth inpainting. In *CVPR*, 2020. [6](#)
- [66] Nikolai Smolyanskiy, Alexey Kamenev, and Stan Birchfield. On the importance of stereo for accurate depth estimation: An efficient semi-supervised deep neural network approach. In *CVPR*, 2018. [2](#)
- [67] Noah Stier, Alexander Rich, Pradeep Sen, and Tobias Höllerer. VoRTX: Volumetric 3D reconstruction with transformers for voxelwise view selection and fusion. In *3DV*, 2021. [8](#)
- [68] Jiaming Sun, Yiming Xie, Linghao Chen, Xiaowei Zhou, and Hujun Bao. NeuralRecon: Real-time coherent 3D reconstruction from monocular video. In *CVPR*, 2021. [8](#)
- [69] Zachary Teed and Jia Deng. DeepV2D: Video to depth with differentiable structure from motion. In *ICLR*, 2020. [2, 6](#)
- [70] Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, and Thomas Brox. DeMoN: Depth and motion network for learning monocular stereo. In *CVPR*, 2017. [6](#)
- [71] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. PatchmatchNet: Learned multi-view patchmatch stereo. In *CVPR*, 2021. [2, 6](#)
- [72] Fangjinhua Wang, Qingtian Zhu, Di Chang, Quankai Gao, Junlin Han, Tong Zhang, Richard Hartley, and Marc Pollefeys. Learning-based multi-view stereo: A survey. *arXiv:2408.15235*, 2024. [2](#)
- [73] Hengyi Wang and Lourdes Agapito. 3D reconstruction with spatial memory. In *3DV*, 2025. [8](#)
- [74] Kaixuan Wang and Shaojie Shen. MVDepthNet: Real-time multiview depth estimation neural network. In *3DV*, 2018. [2, 3](#)
- [75] Rui Wang, Stephen M Pizer, and Jan-Michael Frahm. Recurrent neural network for (un-)supervised learning of monocular video visual odometry and depth. In *CVPR*, 2019. [2](#)
- [76] 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. *arXiv:2410.19115*, 2024. [1, 3](#)
- [77] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. DUST3R: Geometric 3D vision made easy. In *CVPR*, 2024. [2, 5](#)
- [78] 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](#)
- [79] Xiaofeng Wang, Zheng Zhu, Guan Huang, Fangbo Qin, Yun Ye, Yijia He, Xu Chi, and Xingang Wang. MVSTER: Epipolar transformer for efficient multi-view stereo. In *ECCV*, 2022. [3](#)
- [80] Jamie Watson, Oisín Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The Temporal Opportunist: Self-Supervised Multi-Frame Monocular Depth. In *CVPR*, 2021. [2, 3](#)
- [81] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. FoundationStereo: Zero-shot stereo matching. In *CVPR*, 2025. [2](#)
- [82] Juyang Weng, Paul Cohen, Marc Herniou, et al. Camera calibration with distortion models and accuracy evaluation. *PAMI*, 1992. [7](#)
- [83] Felix Wimbauer, Nan Yang, Lukas Von Stumberg, Niclas Zeller, and Daniel Cremers. MonoRec: Semi-supervised dense reconstruction in dynamic environments from a single moving camera. In *CVPR*, 2021. [2](#)
- [84] Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. DepthSplat: Connecting gaussian splatting and depth. In *CVPR*, 2025. [3](#)
- [85] 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. In *ICLR*, 2025. [2, 8](#)
- [86] Jiayu Yang, Jose M Alvarez, and Miaomiao Liu. Self-supervised learning of depth inference for multi-view stereo. In *CVPR*, 2021. [2](#)
- [87] 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. [1, 3](#)
- [88] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth Anything V2. In *NeurIPS*, 2024. [1, 3, 4, 5, 6, 7](#)
- [89] Zhenpei Yang, Zhile Ren, Qi Shan, and Qixing Huang. MVS2D: Efficient multi-view stereo via attention-driven 2D convolutions. In *CVPR*, 2022. [1, 2, 6](#)
- [90] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. MVSNet: Depth inference for unstructured multi-view stereo. In *ECCV*, 2018. [1, 2, 6](#)
- [91] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. BlendedMVS: A large-scale dataset for generalized multi-view stereo networks. In *CVPR*, 2020. [5](#)
- [92] 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. [1, 3](#)
- [93] Zehao Yu and Shenghua Gao. Fast-MVSNet: Sparse-to-dense multi-view stereo with learned propagation and gaussian refinement. In *CVPR*, 2020. [2, 6](#)
- [94] Jure Zbontar and Yann LeCun. Computing the stereo matching cost with a convolutional neural network. In *CVPR*, 2015. [2](#)

- [95] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. GA-Net: Guided aggregation net for end-to-end stereo matching. In *CVPR*, 2019. [2](#)
- [96] Feihu Zhang, Xiaojuan Qi, Ruigang Yang, Victor Prisacariu, Benjamin Wah, and Philip Torr. Domain-invariant stereo matching networks. In *ECCV*, 2020. [2](#)
- [97] Jingyang Zhang, Yao Yao, Shiwei Li, Zixin Luo, and Tian Fang. Visibility-aware multi-view stereo network. In *BMVC*, 2020. [2](#), [6](#)
- [98] Jiawei Zhang, Jiahe Li, Lei Huang, Xiaohan Yu, Lin Gu, Jin Zheng, and Xiao Bai. Robust synthetic-to-real transfer for stereo matching. In *CVPR*, 2024. [2](#)
- [99] Song Zhang, Wenjia Xu, Zhiwei Wei, Lili Zhang, Yang Wang, and Junyi Liu. ARAI-MVSNet: A multi-view stereo depth estimation network with adaptive depth range and depth interval. *Pattern Recognition*, 2023. [3](#), [5](#)