

# Crafting Monocular Cues and Velocity Guidance for Self-Supervised Multi-Frame Depth Learning

Xiaofeng Wang<sup>1,2</sup>, Zheng Zhu<sup>3</sup>, Guan Huang<sup>3</sup>, Xu Chi<sup>3</sup>, Yun Ye<sup>3</sup>, Ziwei Chen<sup>4</sup>, Xingang Wang<sup>1\*</sup>

<sup>1</sup> Institute of Automation, Chinese Academy of Sciences

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Science

<sup>3</sup> PhiGent Robotics

<sup>4</sup> Southeast University

{wangxiaofeng2020, xingang.wang}@ia.ac.cn, zhengzhu@ieee.org  
{guan.huang, xu.chi, yun.ye}@phigent.ai, richard\_chen@seu.edu.cn

## Abstract

Self-supervised monocular methods can efficiently learn depth information of weakly textured surfaces or reflective objects. However, the depth accuracy is limited due to the inherent ambiguity in monocular geometric modeling. In contrast, multi-frame depth estimation methods improve the depth accuracy thanks to the success of Multi-View Stereo (MVS), which directly makes use of geometric constraints. Unfortunately, MVS often suffers from texture-less regions, non-Lambertian surfaces, and moving objects, especially in real-world video sequences without known camera motion and depth supervision. Therefore, we propose MOVEDepth, which exploits the **MO**nocular cues and **VE**locity guidance to improve multi-frame **Depth** learning. Unlike existing methods that enforce consistency between MVS depth and monocular depth, MOVEDepth boosts multi-frame depth learning by directly addressing the inherent problems of MVS. The key of our approach is to utilize monocular depth as a geometric priority to construct MVS cost volume, and adjust depth candidates of cost volume under the guidance of predicted camera velocity. We further fuse monocular depth and MVS depth by learning uncertainty in the cost volume, which results in a robust depth estimation against ambiguity in multi-view geometry. Extensive experiments show MOVEDepth achieves state-of-the-art performance: Compared with Monodepth2 and PackNet, our method relatively improves the depth accuracy by 20% and 19.8% on the KITTI benchmark. MOVEDepth also generalizes to the more challenging DDAD benchmark, relatively outperforming ManyDepth by 7.2%. The code is available at <https://github.com/JeffWang987/MOVEDepth>.

## Introduction

Depth estimation is a fundamental task in 3D computer vision, with versatile applications ranging from virtual/augmented reality (Luo et al. 2020) to autonomous driving (Geiger, Lenz, and Urtasun 2012). Although 3D sensors (e.g., LiDAR, structured light) can generate accurate depth information, it is more attractive to infer depth from a single RGB image in a self-supervised way (Garg et al. 2016;

Godard, Aodha, and Brostow 2016; Godard et al. 2018), which eliminates the necessity of expensive 3D hardware and multi-sensor calibration. However, the accuracy of these monocular methods is not yet on par with 3D sensors due to their inherent ambiguity in geometric modeling.

To improve the monocular depth accuracy, recent multi-frame methods<sup>1</sup> (Watson et al. 2021; Wimbauer et al. 2020; Wang, Pizer, and Frahm 2019; Zhang et al. 2019; Feng et al. 2022; Patil et al. 2020; Wang, Pang, and Lin 2022) leverage temporal and spatially associations in multi-frame video sequences which are available in real-world scenes (e.g., smart devices (Ha et al. 2016) or moving vehicles (Menze and Geiger 2015)). Among these approaches, cost-volume-based methods (Feng et al. 2022; Watson et al. 2021; Wimbauer et al. 2020) achieve state-of-the-art depth accuracy as they take advantage of the successful Multi-View Stereo (MVS). However, MVS is still challenged by unsatisfactory reconstructions in real-world scenes with non-Lambertian surfaces, textureless areas, and moving objects (Knapitsch et al. 2017; Schöps et al. 2017). To tackle these problems, teacher-student training architectures (Watson et al. 2021; Feng et al. 2022; Shao et al. 2022) are proposed to enforce consistency between monocular depth and MVS depth. However, the consistency pushes MVS depth to mimic monocular depth, which underuses the multi-view geometry, thus the performance of these methods is limited.

To improve the multi-view depth accuracy, the learning-based MVS methods (Yao et al. 2018, 2019) densely sample depth candidates in a large range. However, the dense sampling strategy causes matching ambiguity in real-world video frames without known camera pose and depth supervision (see Fig. 3). To mitigate the problem, we explore an efficient approach to enhance geometric cues for improving self-supervised multi-frame depth learning. Our intuition is that the monocular depth serves as a geometric priority of the scene, and the multi-frame matching ambiguity can be significantly reduced by sampling depth candidates near the monocular priority. Apart from the matching ambiguity, the multi-view geometry is still challenged by insufficient *Tri-*

\*Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>Compared with monocular methods that use single frame for inference, the multi-frame methods input  $N$  ( $N \geq 2$ ) frames at inference time.

*angulation Prior* (Schonberger et al. 2016), especially in static/slow video sequences where nearby frames share little stereo baseline. To address the problem, the predicted camera velocity is leveraged to adaptively adjust the depth range. Specifically, multi-frames with higher motion velocity have larger viewpoint change, which can benefit multi-view geometry, so the depth range is enlarged to infer more accurate depth. In contrast, static frames carry little information for depth inference, thus the depth range is shrunk to the more reliable monocular priority. Besides, we fuse monocular depth and MVS depth by learning uncertainty in the cost volume, resulting in a robust depth estimation against artifacts in multi-view geometry (e.g., moving objects, textureless areas).

Owing to the monocular depth priority and camera velocity guidance, MOVEDepth achieves state-of-the-art performance: Compared with competitive monocular baselines (Godard et al. 2018; Guizilini et al. 2019), our method relatively improves the depth accuracy by  $\sim 20\%$  on the KITTI benchmark. MOVEDepth also generalizes to the more challenging DDAD benchmark, relatively outperforming ManyDepth (Watson et al. 2021) by 7.2%. Besides, qualitative analysis demonstrates that our method is more robust against challenging artifacts where multi-view geometry fails.

The main contributions are three-fold as follows:

- We propose a novel self-supervised multi-frame depth learning framework, named MOVEDepth. It leverages the monocular depth cues as a geometric priority, and the multi-frame matching ambiguity is mitigated by sampling depth candidates near the monocular priority.

- The velocity-guided depth sampling is proposed to address failure cases caused by slow/static camera motion. And an adaptive fusing layer is introduced to learn uncertainty in cost volume, which mitigates artifacts brought by textureless areas and moving objects.

- We conduct extensive experiments on KITTI and DDAD, and the results show our method achieves superior depth accuracy in the complex real-world scenes with fewer depth candidates.

## Related Work

In this section, we review depth estimation approaches relevant to our method in the following two categories: monocular depth learning and multi-frame depth learning.

### Monocular Depth Learning

Despite the inherently ill-posed problems of geometric reasoning, monocular depth learning has been studied extensively in the literature. Supervised methods exploit sparse annotations (Chen et al. 2016) or dense points clouds from 3D sensors (Eigen, Puhrsch, and Fergus 2014; Eigen and Fergus 2014; Fu et al. 2018). Self-supervised approaches mitigate the expensive 3D sensors and human annotations, instead training with photo-metric consistency losses using stereo images (Xie, Girshick, and Farhadi 2016; Garg et al. 2016). Recently, monocular video supervision (Zhou et al. 2017; Godard et al. 2018; Guizilini et al. 2019) is an attractive alternative to stereo-based methods but it brings a new

set of challenges: The model needs to jointly estimate camera ego-motion and image depth to form reprojection losses, which further increases the difficulty of monocular geometric modeling. Although there still exists a huge performance gap between monocular methods and other geometric methods, recent advances in monocular depth learning provide a plausible geometric priority, which can benefit other geometric approaches to refine the depth.

### Multi-Frame Depth Learning

Given multiple 2D images and camera parameters, MVS can reconstruct the dense geometry of the scene (Schonberger et al. 2016). Recent advances in learning-based MVS approaches (Yao et al. 2018; Gu et al. 2020; Wang et al. 2022b) further improve the reconstruction accuracy. However, these methods assume known camera poses and static scenes. MonoRec (Wimbauer et al. 2020) extends MVS in dynamic environments that are captured with a single moving camera, but it needs the pre-trained segmentation network (He et al. 2017) and sparse depth obtained by a visual odometry system (Yang et al. 2018). MaGNet (Bae, Budvytis, and Cipolla 2022) fuses single-view probability with multi-view geometry to produce robust estimations. However, it still requires ground truth depth supervision. The self-supervised method (Watson et al. 2021) removes the requirement for human annotations and depth supervision, and it proposes a teacher-student training architecture to encourage the network to ignore unreliable regions in MVS cost volume. (Feng et al. 2022) further improves the depth accuracy in regions of dynamic objects by disentangling object motions, and (Shao et al. 2022) utilizes Deformable Convolution Networks (DCNs (Dai et al. 2017)) to enhance the depth estimates in low-texture and homogeneous-texture regions. However, these methods enforce consistency between MVS depth and monocular depth, which underuses the geometric reasoning of MVS. In contrast, the proposed MOVEDepth directly addresses the inherent problems of MVS (e.g., static/slow camera motion, object motions, and textureless regions).

## Method

A detailed description of MOVEDepth is given in this section and the main network architecture is illustrated in Fig. 1. Given a video sequence, (a) we firstly utilize PoseNet to estimate camera ego-motion and velocity between frame  $T$  and frame  $T - 1$ . (b) Then the monocular depth is predicted using the DepthNet. (c) Subsequently, we conduct homography warping between the encoded frame features using the predicted camera ego-motion and monocular depth priority. The resulting cost volume is decoded into a depth map and an uncertainty map, which serves as guidance for fusing the monocular depth and MVS depth. (d) Significantly, the depth sampling range of homography warping is adaptively adjusted under the guidance of predicted camera velocity, which mitigates problems brought by slow/static camera motion.

The following parts start with the preliminary of self-supervised monocular depth learning from a video sequence.

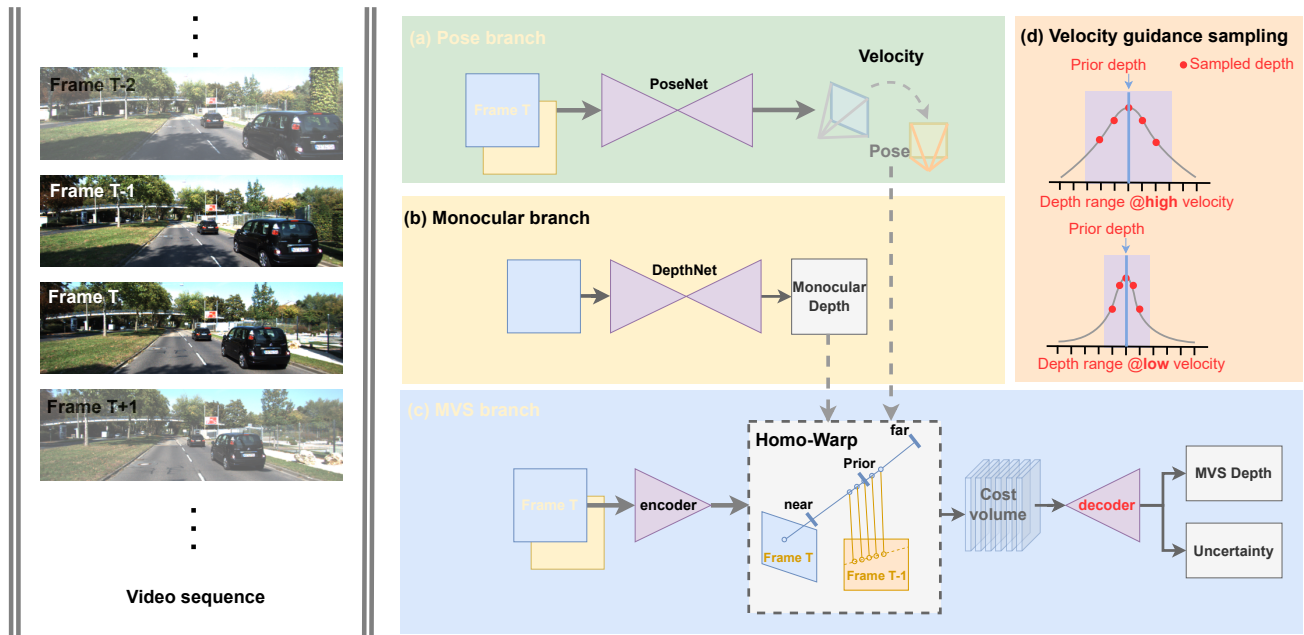


Figure 1: The main network architecture of MOVEDepth. (a) PoseNet is utilized to estimate camera ego-motion and velocity between frame  $T$  and frame  $T - 1$ . (b) The monocular depth is predicted using the DepthNet, which serves as a geometric priority to construct MVS cost volume. (c) We conduct homography warping between the encoded frame features using the predicted camera ego-motion and monocular depth priority. The resulting cost volume is decoded into a depth map and an uncertainty map. (d) The depth sampling range is adaptively adjusted under the guidance of predicted camera velocity.

We then introduce the important innovations to improve multi-frame depth learning by fusing monocular cues.

### Self-Supervised Monocular Depth Learning

The self-supervised pipeline is conducted by jointly training a DepthNet  $\theta_d$  (see Fig. 1(b)) and a PoseNet  $\theta_p$  (see Fig. 1(a)) (Zhou et al. 2017), and they are trained only on video frames  $\{\mathbf{I}_t\}_{t=1}^N$ . Specifically, we estimate the monocular depth  $D_{\text{Mono}} = \theta_d(\mathbf{I}_t)$  of current frame  $\mathbf{I}_t$ , and predict relative camera pose  $[\mathbf{R} \mid \mathbf{T}]_{t \rightarrow t+k} = \theta_p(\mathbf{I}_t, \mathbf{I}_{t+k})$  between frame  $\mathbf{I}_t$  and frame  $\mathbf{I}_{t+k}$  ( $k \in \{-1, 1\}$ ). Then, we can synthesize  $\mathbf{I}_t$  from viewpoint  $\mathbf{I}_{t+k}$  by the following operation:

$$\mathbf{I}_{t+k \rightarrow t}(D_{\text{Mono}}) = \mathbf{I}_{t+k} \langle \text{proj}(D_{\text{Mono}}, [\mathbf{R} \mid \mathbf{T}]_{t \rightarrow t+k}, \mathbf{K}) \rangle, \quad (1)$$

where  $\mathbf{K}$  is the camera intrinsics,  $\text{proj}(\cdot)$  is the projection function that returns the 2D pixel coordinates of the projected  $D_{\text{Mono}}$ , and  $\langle \cdot \rangle$  is the pixel sampling operator. Following the optimization convention (Godard et al. 2018), the training pipeline is optimized by a reprojection loss:

$$\mathcal{L}_r(D_{\text{Mono}}) = \min_k pe(\mathbf{I}_t, \mathbf{I}_{t+k \rightarrow t}(D_{\text{Mono}})), \quad (2)$$

where the min operation selects the best matching frames to avoid ambiguity brought by occlusions, and  $pe(\cdot)$  is a weighted combination of  $\mathcal{L}_1$  loss and structure similarity (SSIM) loss. The reprojection loss is calculated over multi-scale depth outputs, and more implementation details can be found in (Godard et al. 2018).

### MOVEDepth Design

**MVS Depth from Monocular Depth Priority** Multi-view approaches warp source images into the reference camera frustum to form cost volume, and estimate depth to be the highest-activated value in cost volume (Yao et al. 2018). Although the hard-coded multi-view methods reduce geometry ambiguity and generate more accurate depth, they are still challenged by texture-less regions, non-Lambertian surfaces, and moving objects, especially in real-world video frames without known camera motion. The monocular methods, on the other hand, are more robust against weakly textured regions or moving objects but the overall depth accuracy is limited. Therefore, we exploit monocular cues to complement the limitation of MVS (see Fig. 1(c)), which is elaborated in the following.

Given a current frame  $\mathbf{I}_t \in \mathbb{R}^{H \times W \times 3}$  and its nearby frame  $\mathbf{I}_{t-1} \in \mathbb{R}^{H \times W \times 3}$  (the future frames is not used to enable online depth prediction), we firstly leverage an encoder  $\theta_{\text{enc}}$  to extract 2D features of these frames, where the images are downscaled to lower resolution deep features  $\mathbf{F}_i (i \in \{0, -1\}) \in \mathbb{R}^{H/4 \times W/4 \times C}$ . Following previous learning-based MVS (Yao et al. 2018; Gu et al. 2020), the plane sweep stereo (Collins 1996) is utilized to establish multiple front-to-parallel planes in the current frame. Specifically, equipped with camera intrinsic  $\mathbf{K}$  and extrinsic  $[\mathbf{R} \mid \mathbf{T}]$  estimated by PoseNet  $\theta_p$ , the previous frame features can be warped into the current camera frustum:

$$\mathbf{p}_{t-1,j} = \mathbf{K} \cdot (\mathbf{R} \cdot (\mathbf{K}^{-1} \cdot \mathbf{p}_t \cdot d_j) + \mathbf{T}), \quad (3)$$

where  $d_j$  is the  $j$ -th hypothesized depth candidates of pixel  $\mathbf{p}_t$  in the current frame feature  $\mathbf{F}_t$ , and  $\mathbf{p}_{t-1,j}$  denotes the corresponding pixel in the previous frame feature  $\mathbf{F}_{t-1}$ . After the warping operation, the volume feature  $\mathbf{V}_{t-1} \in \mathbb{R}^{H/4 \times W/4 \times C \times D}$  is constructed, where  $D$  is the number of depth candidates. Significantly, to reduce depth searching space, we specify the depth range  $\mathcal{R}$  using the monocular depth priority  $D_{\text{Mono}}$ :

$$\mathcal{R} = \{d | d_{\min} \leq d \leq d_{\max}\}, \quad (4)$$

where  $(d_{\min} + d_{\max})/2 = D_{\text{Mono}}$ , and  $d_{\min}, d_{\max}$  is adaptively adjusted under the guidance of camera velocity, which is elaborated in the next subsection.

Given previous frame volume  $\mathbf{V}_{t-1}$ , we then use *group correlation* (Wang et al. 2021; Xu and Tao 2019) to construct cost volume, which measures the visual similarity between the current frame and the previous frame:

$$\mathbf{s}_i^g = \frac{1}{G} \langle \mathbf{v}_i^g, \mathbf{f}_i^g \rangle, \quad (5)$$

where  $\mathbf{v}_i^g \in \mathbb{R}^{C \times D}$  is the  $g$ -th group feature of  $\mathbf{v}_i$  ( $\mathbf{v}_i \in \mathbb{R}^{C \times D}$  is the  $i$ -th pixel feature of  $\mathbf{V}_t$ ), and  $\mathbf{f}_i^g \in \mathbb{R}^{C \times 1}$  is the  $g$ -th group feature of  $\mathbf{f}_i$  ( $\mathbf{f}_i$  is the  $i$ -th pixel feature of  $\mathbf{F}_t$ ), and  $\langle \cdot, \cdot \rangle$  is the inner product. Then  $\{\mathbf{s}_i^g\}_{g=0}^{G-1}$  are channel-wise stacked to generate  $\mathbf{s}_i \in \mathbb{R}^{G \times D}$ , which is  $i$ -th pixel feature of the final cost volume  $\mathbf{S}_t \in \mathbb{R}^{H/4 \times W/4 \times G \times D}$ .

The calculated cost volume is subsequently decoded by a light-weight  $\theta_{\text{dec}}$  to get depth probability  $\mathbf{P} \in \mathbb{R}^{H/4 \times W/4 \times D}$ , and the MVS depth is generated by *localmax* (Wang et al. 2022a):

$$D_{\text{MVS}}(\mathbf{p}) = \left( \frac{1}{\sum_{j=\mathbf{X}(\mathbf{p})-r}^{\mathbf{X}(\mathbf{p})+r} \mathbf{p}_j} \sum_{j=\mathbf{X}(\mathbf{p})-r}^{\mathbf{X}(\mathbf{p})+r} \frac{1}{d_j} \cdot \mathbf{p}_j} \right)^{-1}, \quad (6)$$

where  $\mathbf{p} \in \mathbb{R}^D$  is the pixel value of  $\mathbf{P}$ ,  $\mathbf{X}(\mathbf{p}) = \text{argmax}_j \mathbf{p}_j$  is the index of the highest value for  $\mathbf{p}$ , and  $r$  is a radius parameter (typically set as 1). Finally, the *convex interpolation* (Teed and Deng 2020) is leveraged to upsample the MVS depth to the original resolution.

**Velocity-Guided Depth Sampling** In the previous subsection, the monocular depth is leveraged as a geometric center for depth sampling, but the depth range is left to be addressed. Typically, learning-based MVS (Yao et al. 2018, 2019) sample depth candidates in a fixed range, which is either calculated by COLMAP (Schonberger et al. 2016) or learned by networks (Watson et al. 2021). However, the depth range is utilized to describe the entire scene, and densely searching in such a wide range is computationally expensive and can not produce accurate depth (Gu et al. 2020). Recent methods reduce depth range by coarse-to-fine sampling (Gu et al. 2020; Wang et al. 2021) or confidence-based sampling (Bae, Budvytis, and Cipolla 2022; Cheng et al. 2020). However, we empirically find these sampling strategies are limited in self-supervised multi-frame depth learning (see Tab. 4), as they overlook the *Triangulation Prior* (Schonberger et al. 2016) of nearby frames.




Figure 2: MOVEDDepth learns uncertainty in depth probability to fuse monocular depth and MVS depth. The upper branch decodes depth probability into depth by *localmax* (Eq. 6), and the lower branch decodes depth probability into an uncertainty map, which serves as a guidance for fusing MVS depth and monocular depth.

To mitigate the problem, we propose velocity-guided depth sampling. The key innovation is to associate *Triangulation Prior* with camera motion velocity  $v$ . Namely, the viewpoint changes noticeably when the camera moves at a high velocity, providing a sufficient *Triangulation Prior* for multi-view geometry. In contrast, slow/static video frames share a similar viewpoint, thus the *Triangulation Prior* is limited (theoretical analysis is in supplement). For video frames with sufficient *Triangulation Prior*, we expand depth range to infer accurate depth, and for frames with insufficient *Triangulation Prior*, the depth range is shrunk to the more reliable monocular priority. The depth sampling range is specified as follows:

$$\begin{aligned} d_{\min} &= D_{\text{Mono}}(1 - \beta\mathcal{T}(v)) \\ d_{\max} &= D_{\text{Mono}}(1 + \beta\mathcal{T}(v)), \end{aligned} \quad (7)$$

where the camera motion velocity  $v = \alpha \|\mathbf{T}\|_2$  is the byproduct of PoseNet  $\theta_p$  ( $\mathbf{T}$  is the camera translation estimated by  $\theta_p$ , and  $\alpha$  is the camera frame rate).  $\beta$  is a hyper-parameter, and  $\mathcal{T}(\cdot)$  is a scale function that transforms  $v$  to a real-world scale, which can be calculated by median-scaling (Godard et al. 2018) or camera-height-scaling (Yin et al. 2017). To ensure training stability,  $\beta\mathcal{T}(\cdot)$  is clamped to range (0, 1).

Notably, the depth sampling strategy resembles *Gaussian Sampling* with mean of  $D_{\text{Mono}}$  and variance of  $\beta\mathcal{T}(v)$ , and the MVS depth range shrinks to the more reliable monocular depth when the camera is static. Differently, the depth candidates are not sampled by their probability, but by a deterministic inverse sampling strategy:

$$d_j = \left( \left( \frac{1}{d_{\min}} - \frac{1}{d_{\max}} \right) \frac{j}{D-1} + \frac{1}{d_{\max}} \right)^{-1}, \quad (8)$$

where  $j = 0 \dots D-1$ . Compared with linear sampling (Yao et al. 2018) or probabilistic sampling (Bae, Budvytis, and Cipolla 2022), inverse depth sampling results in uniformly distributed depth candidates at the pixel level, which is beneficial for large-scale multi-frame matching (Xu and Tao 2019).

**Uncertainty-Based Depth Fusing** The calculated  $D_{\text{MVS}}$  is still challenged by texture-less regions, non-Lambertian

surfaces, and moving objects, which are the inherent problems of the multi-view geometry. To alleviate the problem, the uncertainty-based fusing method is introduced to replace unsatisfactory  $D_{MVS}$  with the more reliable  $D_{Mono}$ . As shown in Fig. 2, we leverage an *Uncertainty Decoder*  $\theta_u$  to learn uncertainty map  $\mathbf{U}$  from the entropy of depth probability  $\mathbf{p}$ :

$$\mathbf{U}(\mathbf{p}) = \theta_u \left( \sum_{j=0}^{D-1} -\mathbf{p}_j \log \mathbf{p}_j \right), \quad (9)$$

where  $\theta_u$  is comprised of 2D Convolutional Neural Network (CNN) blocks and a *Sigmoid* function. The reason for adopting the entropy is that the randomness of the depth probability distribution is positively related to the MVS depth uncertainty (Zhang et al. 2020). Subsequently, the uncertainty map is leveraged to calculate the fused depth  $D_{Fuse}$ :

$$D_{Fuse} = \mathbf{U} \odot D_{Mono} + (\mathbf{1} - \mathbf{U}) \odot D_{MVS}, \quad (10)$$

where  $\odot$  denotes element-wise product.

**Loss Function** MOVEDepth is end-to-end trained in a self-supervised manner, and the loss consists of three parts:

$$\mathcal{L}_{MOVEDepth} = \lambda_1 \mathcal{L}(D_{Mono}) + \lambda_2 \mathcal{L}(D_{MVS}) + \lambda_3 \mathcal{L}(D_{Fuse}), \quad (11)$$

where  $\lambda_1, \lambda_2, \lambda_3$  are the loss weights, and  $\mathcal{L}(\cdot)$  is a weighted combination of reprojection loss  $\mathcal{L}_r$  (Eq. (2)) and depth smooth loss  $\mathcal{L}_s$  (Godard, Aodha, and Brostow 2016):

$$\mathcal{L}(D) = \mathcal{L}_r(D) + \gamma \mathcal{L}_s(D), \quad (12)$$

where  $\gamma$  denotes the loss weight.

## Experiment

### Datasets

MOVEDepth is evaluated on KITTI (Geiger, Lenz, and Urtasun 2012) and DDAD (Guizilini et al. 2019) to verify the effectiveness. KITTI is an outdoor dataset in the driving scenario, which is the standard benchmark for depth evaluation. Following from the Eigen split (Eigen and Fergus 2014), with data preprocessing from (Zhou et al. 2017), the data is divided into 39810/4424/697 training, validation and test images. For DDAD, it is a diverse dataset of the highway, urban, and residential scenes curated by self-driving cars, which is a novel benchmark for depth evaluation. Notably, DDAD evaluates with longer depth ranges and denser LiDAR ground-truth, which is particularly challenging for multi-frame methods. Following (Guizilini et al. 2019), only front-view images are used, resulting in 12560/3950 training and validation images.

### Implementation Details

Following previous work (Godard et al. 2018; Watson et al. 2021), we use color-jitter and flip as training-time augmentations, and MOVEDepth is trained with an input resolution of  $640 \times 192$  (KITTI) and  $640 \times 384$  (DDAD). We only use two frames  $\{\mathbf{I}_{t-1}, \mathbf{I}_t\}$  for cost volume construction, and use  $\{\mathbf{I}_{t-1}, \mathbf{I}_t, \mathbf{I}_{t+1}\}$  for reprojection loss. We train MOVEDepth

for 20 epochs and optimize it with Adam (Kingma and Ba 2015). MOVEDepth is trained on 4 NVIDIA RTX 3090 GPUs with batch size 6 on each GPU. The learning rate is initially set as 0.0002, which decays by a factor of 10 for the final 5 epochs. The feature extractor in  $\theta_d$  and  $\theta_p$  follows the same architecture in (Godard et al. 2018; Guizilini et al. 2019), and  $\theta_{enc}, \theta_{dec}$  comprise light-weight 2D CNNs and 3D CNNs (more network details are in supplement). Following (Godard et al. 2018; Watson et al. 2021), the loss weight  $\gamma$  is set as 0.001, and  $\lambda_{i(i \in \{1,2,3\})} = 1$ . For MVS cost volume construction, the number of depth candidates is 16, group correlation  $G = 16$ , and  $\beta = 0.15$ .

### KITTI Results

We compare MOVEDepth with monocular methods and multi-frame methods, and the quantitative results are shown in Tab. 1, where MOVEDepth achieves a state-of-the-art performance among all the competitors. We observe that MOVEDepth relatively improves the best-performing monocular method (Guizilini et al. 2020) by 12.7% (on Abs. Rel.), which demonstrates it is significant to investigate multi-frame geometry to improve depth accuracy. Besides, MOVEDepth relatively outperforms multi-frame methods (ManyDepth (Watson et al. 2021) and DynamicDepth (Feng et al. 2022)) by 9.2% and 7.3% (on Abs. Rel.), which verifies that our method is effective for multi-frame depth learning. Furthermore, we implement MOVEDepth with depth priority from different monocular methods, and our method relatively improves Monodepth2 (Godard et al. 2018) and PackNet (Guizilini et al. 2019) by 20% and 19.8% (on Abs. Rel.), which verifies that MOVEDepth benefits from advances in existing monocular methods. Notably, the resolution of depth priority is  $\frac{1}{4}$  of the inputs, and the depth priority significantly reduces the depth candidates of MVS, thus MOVEDepth (with Monodepth2) still runs at an acceptable speed (32 FPS) at single RTX 3090.

Qualitative results are presented in Fig. 3, where the MVS baseline (without monocular priority) fails in video-based depth learning due to matching ambiguity in a large depth range. In contrast, monocular method (Godard et al. 2018) and MOVEDepth can generate plausible depth prediction. We further visualize the absolute relative depth error, which shows that MOVEDepth can significantly reduce the depth error brought by monocular methods.

### DDAD Results

DDAD is a novel benchmark for depth estimation, which evaluates depth with longer distance and denser LiDAR points. We conduct experiments on DDAD to verify our method can generalize to more challenging scenarios. As shown in Tab. 2, MOVEDepth still outperforms previous competitors. Notably, our method relatively improves the state-of-the-art multi-frame method (Watson et al. 2021) by 7.6% (on Abs. Rel.).

### Ablation Study

In this subsection, we conduct an ablation study on the KITTI dataset to analyze the effectiveness of each component. Firstly, we analyze the impact of the monocular

| Method                                  | Test frames   | W × H      | Abs. Rel.    | Sq. Rel.     | RMSE         | RMSE <sub>log</sub> | $\delta < 1.25$ |
|---|---------------|------------|--------------|--------------|--------------|---------------------|-----------------|
| EPC++ (Luo et al. 2019)                 | 1             | 832 × 256  | 0.141        | 1.029        | 5.350        | 0.216               | 0.816           |
| Struct2depth (M) (Casser et al. 2019)   | 1             | 416 × 128  | 0.141        | 1.026        | 5.291        | 0.215               | 0.816           |
| Videos in the wild (Gordon et al. 2019) | 1             | 416 × 128  | 0.128        | 0.959        | 5.230        | 0.212               | 0.845           |
| Monodepth2 (Godard et al. 2018)         | 1             | 640 × 192  | 0.115        | 0.903        | 4.863        | 0.193               | 0.877           |
| PackNet-SFM (Guizilini et al. 2019)     | 1             | 640 × 192  | 0.111        | 0.785        | 4.601        | 0.189               | 0.878           |
| R-MSFM (Zhou et al. 2021)               | 1             | 1024 × 320 | 0.108        | 0.748        | 4.481        | 0.179               | 0.893           |
| RM-Depth (Hui 2022)                     | 1             | 640 × 192  | 0.108        | 0.710        | 4.513        | 0.183               | 0.884           |
| GLNet (Chen et al. 2019)                | 3 (-1, 0, +1) | 416 × 128  | 0.099        | 0.796        | 4.743        | 0.186               | 0.884           |
| CoMoDA (Kuznietsov et al. 2021)         | N             | 640 × 192  | 0.103        | 0.862        | 4.594        | 0.183               | 0.899           |
| Patil (Patil et al. 2020)               | N             | 640 × 192  | 0.111        | 0.821        | 4.650        | 0.187               | 0.883           |
| TC-Depth (Ruhkamp et al. 2021)          | 3(-1, 0, +1)  | 640 × 192  | 0.103        | 0.746        | 4.483        | 0.185               | 0.894           |
| ManyDepth (Watson et al. 2021)          | 2 (-1, 0)     | 640 × 192  | 0.098        | 0.770        | 4.459        | 0.176               | 0.900           |
| DynamicDepth (Feng et al. 2022)         | 2 (-1, 0)     | 640 × 192  | 0.096        | 0.720        | 4.458        | 0.175               | 0.897           |
| MOVEDepth (with Monodepth2)             | 2 (-1, 0)     | 640 × 192  | 0.092        | 0.686        | 4.332        | 0.173               | <b>0.904</b>    |
| <b>MOVEDepth (with PackNet)</b>         | 2 (-1, 0)     | 640 × 192  | <b>0.089</b> | <b>0.663</b> | <b>4.216</b> | <b>0.169</b>        | <b>0.904</b>    |

Table 1: Comparison of MOVEDepth to existing self-supervised methods on the KITTI (Geiger, Lenz, and Urtasun 2012) Eigen split. At top we compare MOVEDepth with monocular methods using one frame at test time. At middle we compare MOVEDepth with multi-frame methods with multiple frames inputs at test time. At bottom we show our method with different monocular depth priority.

| Method                   | Abs. Rel.    | Sq. Rel.     | $\delta < 1.25$ |
|--------------------------|--------------|--------------|-----------------|
| Monodepth2               | 0.213        | 4.975        | 0.761           |
| Packnet                  | 0.162        | 3.917        | 0.823           |
| GUDA                     | 0.147        | 2.922        | 0.809           |
| ManyDepth                | 0.145        | 3.246        | 0.821           |
| Ours (w/ Monodepth2)     | 0.136        | 3.027        | 0.835           |
| <b>Ours (w/ Packnet)</b> | <b>0.134</b> | <b>2.903</b> | <b>0.837</b>    |

Table 2: Depth evaluation results on DDAD validation set (Guizilini et al. 2019). The above methods are evaluated for depth range of 200m without cropping.

priority for MVS, demonstrating that the monocular priority significantly reduces multi-frame matching ambiguity with fewer depth candidates. Then we verify that, compared with cascade sampling (Gu et al. 2020) and confidence sampling (Bae, Budvytis, and Cipolla 2022), multi-frame depth learning benefits more from the proposed velocity-guided depth sampling. Lastly, we qualitatively show that the uncertainty-based fusing makes MOVEDepth more robust against challenging artifacts where multi-view geometry fails (e.g., moving objects). In the following ablation study, unless specified, we utilize Monodepth2 (Godard et al. 2018) as the monocular priority, and only two frames (0, -1) are leveraged to construct MVS cost volume.

### Monocular Priority for Multi-Frame Depth Learning

As shown in Fig. 3, due to matching ambiguity in a large depth range, the multi-view geometry (without monocular priority) fails to learn multi-frame depth in a self-supervised manner. Besides, quantitative results (Tab. 3) also show that the MOVEDepth (without monocular priority) is inferior to the monocular priority (Godard et al. 2018), yet with  $3 \times \sim 4 \times$  GPU memory consumption.

| Method                   | Abs. Rel.    | $\delta < 1.25$ | GPU (MB)   |
|--------------------------|--------------|-----------------|------------|
| Mono. baseline           | 0.115        | 0.877           | <b>161</b> |
| Ous w/o prior. (96 bins) | 0.328        | 0.633           | 678        |
| Ous w/o prior. (48 bins) | 0.349        | 0.582           | 483        |
| Ous (48 bins)            | 0.096        | 0.900           | 466        |
| <b>Ous (16 bins)</b>     | <b>0.092</b> | <b>0.904</b>    | 322        |
| Ous (8 bins)             | 0.094        | 0.903           | 285        |

Table 3: Ablation analysis on monocular priority for multi-frame depth learning, including depth metrics and inference GPU memory consumption (GPU consumption is measured by `torch.cuda.max_memory_allocated`).

When integrated with the monocular depth priority, the geometric uncertainty is significantly reduced, and MOVEDepth can produce superior depth predictions with less depth bin candidates. We empirically select 16 depth bins for MOVEDepth, as it strikes a balance between accuracy and memory consumption, and MOVEDepth with 16 depth bins relatively outperform the monocular baseline by 20% (on Abs. Rel.).

**Velocity-Guided Depth Sampling** As shown in Tab. 4, MOVEDepth with fixed depth sampling range (row 1,2) shows restricted performance, and increasing input frames (row 3,4) can not improve depth accuracy due to the inherent geometric ambiguity (without pose and depth supervision). Recent methods adjust depth range by cascade sampling (Gu et al. 2020) (half the depth range in the next stage) or confidence-based sampling (Bae, Budvytis, and Cipolla 2022) (adjust depth range according to monocular depth confidence). However, these sampling strategies (row 5,6) bring marginal depth accuracy improvement, as they overlook the *Triangulation Prior* nearby frames. In contrast,

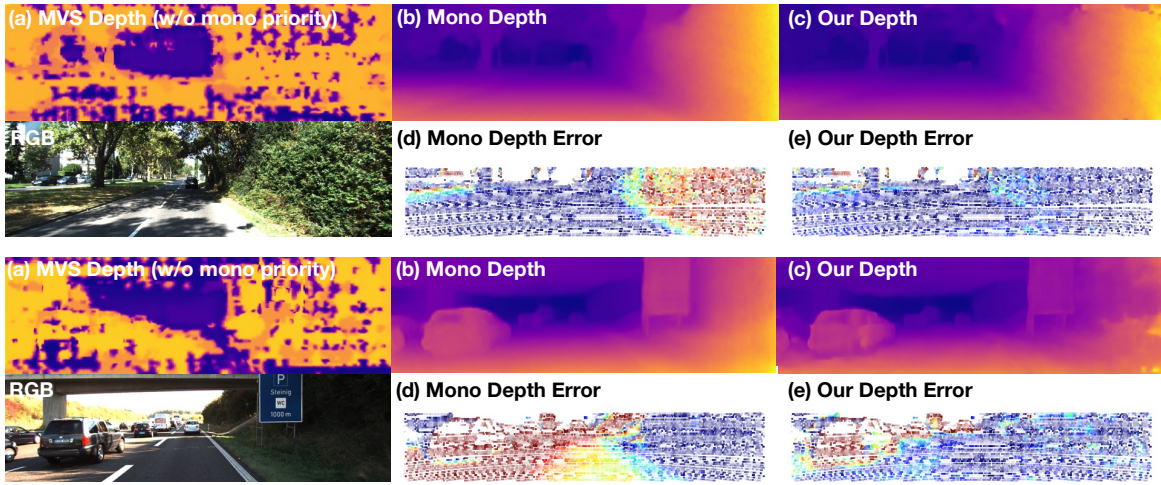


Figure 3: Qualitative results on KITTI. We visualize depth prediction from (a) MVS baseline (MOVEDepth without monocular priority), (b) monocular method (Godard et al. 2018), and (c) MOVEDepth. (d)(e) We also visualize the absolute relative error (Abs. Rel.) compared to the ground truth, where the error ranges from blue (Abs. Rel. = 0.0) to red (Abs. Rel. = 0.2).

| Method                                     | Abs. Rel.    | Sq. Rel.     | $\delta < 1.25$ |
|--|--------------|--------------|-----------------|
| Ours <sup>†</sup> (frame 0, -1)            | 0.101        | 0.801        | 0.897           |
| Ours <sup>‡</sup> (frame 0, -1)            | 0.099        | 0.773        | 0.898           |
| Ours <sup>†</sup> (frame 0, -1, +1)        | 0.098        | 0.769        | 0.900           |
| Ours (cascade)                             | 0.096        | 0.762        | 0.899           |
| Ours (confidence)                          | 0.097        | 0.781        | 0.897           |
| Ours (vel. $\beta=0.1$ )                   | 0.094        | 0.712        | 0.902           |
| <b>Ours (vel. <math>\beta=0.15</math>)</b> | <b>0.092</b> | <b>0.686</b> | <b>0.904</b>    |
| Ours (vel. $\beta=0.2$ )                   | 0.093        | 0.692        | 0.902           |

Table 4: Ablation analysis on velocity-guided depth sampling. MOVEDepth<sup>†</sup> denotes fixed depth range  $[\frac{1}{2}D_{\text{Mono}}, \frac{3}{2}D_{\text{Mono}}]$ , and MOVEDepth<sup>‡</sup> denotes fixed depth range  $[\frac{3}{4}D_{\text{Mono}}, \frac{5}{4}D_{\text{Mono}}]$ .

| Method         | Abs. Rel.    | Sq. Rel.     | RMSE         | $\delta < 1.25$ |
|----------------|--------------|--------------|--------------|-----------------|
| Ours (unfused) | 0.094        | 0.728        | 4.362        | 0.902           |
| Ours (fused)   | <b>0.092</b> | <b>0.686</b> | <b>4.332</b> | <b>0.904</b>    |

Table 5: Ablation analysis on uncertainty-based fusing.

MOVEDepth associates *Triangulation Prior* with the predicted camera velocity, and adaptively adjusts depth range under the guidance of velocity. Significantly, MOVEDepth with velocity guidance ( $\beta=0.15$ ) relatively outperforms its fixed depth range baseline (row 1) by 7.1% (on Abs. Rel.).

**Uncertainty-Based Depth Fusing** The uncertainty-based depth fusing strategy is utilized to complement challenging artifacts where multi-view geometry fails. As shown in Fig. 4, the moving objects and textureless areas are learned as highly-uncertain regions. Besides, learning the uncertainty in MVS regularizes the cost volume and further improves the depth accuracy on all the metrics (see Tab. 5).



Figure 4: Visualization of the uncertainty map, where the left figures are RGB images, and the right figures are the learned uncertainty maps (white: certain, black: uncertain).

## Conclusion

In this paper, we propose MOVEDepth, which crafts monocular cues and velocity guidance for improving multi-frame depth learning in a self-supervised manner. The proposed method leverages monocular depth priority to reduce multi-frame matching geometry ambiguity. Specifically, MOVEDepth constructs cost volume with depth range centered at the monocular priority, and the depth range is adaptively adjusted by the predicted camera velocity. Therefore, video frames with insufficient *Triangulation Prior* shrink the depth range to monocular priority, and those with sufficient *Triangulation Prior* expand the depth range to infer more accurate depth. Moreover, we learn depth uncertainty from the entropy of MVS depth probability, which results in a more robust depth against challenging artifacts where multi-view geometry fails. Extensive experiments on KITTI and DDAD show that MOVEDepth achieves state-of-the-art depth accuracy with significantly reduced depth range. We hope that MOVEDepth can inspire more self-supervised multi-frame depth learning methods in the future.

## References

- Bae, G.; Budvytis, I.; and Cipolla, R. 2022. Multi-View Depth Estimation by Fusing Single-View Depth Probability with Multi-View Geometry. In *CVPR*.
- Casser, V.; Pirk, S.; Mahjourian, R.; and Angelova, A. 2019. Depth Prediction Without the Sensors: Leveraging Structure for Unsupervised Learning from Monocular Videos. In *AAAI*.
- Chen, W.; Fu, Z.; Yang, D.; and Deng, J. 2016. Single-Image Depth Perception in the Wild. In *NeurIPS*.
- Chen, Y.; Schmid, C.; Sminchisescu; and Cristian. 2019. Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In *ICCV*.
- Cheng, S.; Xu, Z.; Zhu, S.; Li, Z.; Li, L. E.; Ramamoorthi, R.; and Su, H. 2020. Deep Stereo Using Adaptive Thin Volume Representation With Uncertainty Awareness. In *CVPR*.
- Collins, R. T. 1996. A space-sweep approach to true multi-image matching. In *CVPR*.
- Dai, J.; Qi, H.; Xiong, Y.; Li, Y.; Zhang, G.; Hu, H.; and Wei, Y. 2017. Deformable Convolutional Networks. In *ICCV*.
- Eigen, D.; and Fergus, R. 2014. Predicting Depth, Surface Normals and Semantic Labels with a Common Multi-Scale Convolutional Architecture. In *ICCV*.
- Eigen, D.; Puhersch, C.; and Fergus, R. 2014. Depth Map Prediction from a Single Image using a Multi-Scale Deep Network. In *NeurIPS*.
- Feng, Z.; Yang, L.; Jing, L.; Wang, H.; Tian, Y.; and Li, B. 2022. Disentangling Object Motion and Occlusion for Unsupervised Multi-frame Monocular Depth. In *ECCV*.
- Fu, H.; Gong, M.; Wang, C.; Batmanghelich, K.; and Tao, D. 2018. Deep Ordinal Regression Network for Monocular Depth Estimation. In *CVPR*.
- Garg, R.; Kumar, B. G. V.; Carneiro, G.; and Reid, I. 2016. Unsupervised CNN for Single View Depth Estimation: Geometry to the Rescue. In *ECCV*.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *CVPR*.
- Godard, C.; Aodha, O. M.; and Brostow, G. J. 2016. Unsupervised Monocular Depth Estimation with Left-Right Consistency. In *CVPR*.
- Godard, C.; Aodha, O. M.; Firman, M.; and Brostow, G. J. 2018. Digging Into Self-Supervised Monocular Depth Estimation. In *ICCV*.
- Gordon, A.; Li, H.; Jonschkowski, R.; and Angelova, A. 2019. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *ICCV*.
- Gu, X.; Fan, Z.; Zhu, S.; Dai, Z.; Tan, F.; and Tan, P. 2020. Cascade Cost Volume for High-Resolution Multi-View Stereo and Stereo Matching. In *CVPR*.
- Guizilini, V.; Ambrus, R.; Pillai, S.; Raventos, A.; and Gaidon, A. 2019. 3D Packing for Self-Supervised Monocular Depth Estimation. In *CVPR*.
- Guizilini, V.; Hou, R.; Li, J.; Ambrus, R.; and Gaidon, A. 2020. Semantically-Guided Representation Learning for Self-Supervised Monocular Depth. In *ICLR*.
- Ha, H.; Im, S.; Park, J.; Jeon, H.-G.; and Kweon, I. S. 2016. High-Quality Depth from Uncalibrated Small Motion Clip. In *CVPR*.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. B. 2017. Mask R-CNN. In *ICCV*.
- Hui, T.-W. 2022. RM-Depth: Unsupervised Learning of Recurrent Monocular Depth in Dynamic Scenes. In *CVPR*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- Knapitsch, A.; Park, J.; Zhou, Q.; and Koltun, V. 2017. Tanks and temples: benchmarking large-scale scene reconstruction. *ToG*.
- Kuznetsov; Yevhen; Proesmans; Marc; and Van Gool, L. 2021. CoMoDA: Continuous Monocular Depth Adaptation Using Past Experiences. In *WACV*.
- Luo, C.; Yang, Z.; Wang, P.; Wang, Y.; Xu, W.; Nevatia, R.; and Yuille, A. 2019. Every pixel counts++: Joint learning of geometry and motion with 3D holistic understanding. *PAMI*.
- Luo, X.; Huang, J.-B.; Szeliski, R.; Matzen, K.; and Kopf, J. 2020. Consistent video depth estimation. In *SIGGRAPH*.
- Menze, M.; and Geiger, A. 2015. Object scene flow for autonomous vehicles. In *CVPR*.
- Patil, V.; Van Gansbeke, W.; Dai, D.; and Van Gool, L. 2020. Don't Forget The Past: Recurrent Depth Estimation from Monocular Video. In *RA-L*.
- Ruhkamp, P.; Gao, D.; Chen, H.; Navab, N.; and Busam, B. 2021. Attention meets Geometry: Geometry Guided Spatial-Temporal Attention for Consistent Self-Supervised Monocular Depth Estimation. In *3DV*.
- Schonberger, J. L.; Zheng, E.; Frahm, J.-M.; and Pollefeys, M. 2016. Pixelwise View Selection for Unstructured Multi-View Stereo. In *ECCV*.
- Schöps, T.; Schönberger, J. L.; Galliani, S.; Sattler, T.; Schindler, K.; Pollefeys, M.; and Geiger, A. 2017. A Multi-view Stereo Benchmark with High-Resolution Images and Multi-camera Videos. In *CVPR*.
- Shao, S.; Pei, Z.; Chen, W.; Wu, X.; Liu, Z.; and Li, Z. 2022. SMUDLP: Self-Teaching Multi-Frame Unsupervised Endoscopic Depth Estimation with Learnable Patchmatch. arXiv:2205.15034.
- Teed, Z.; and Deng, J. 2020. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *ECCV*.
- Wang, F.; Galliani, S.; Vogel, C.; and Pollefeys, M. 2022a. IterMVS: Iterative Probability Estimation for Efficient Multi-View Stereo. In *CVPR*.
- Wang, F.; Galliani, S.; Vogel, C.; Speciale, P.; and Pollefeys, M. 2021. PatchmatchNet: Learned Multi-View Patchmatch Stereo. In *CVPR*.
- Wang, R.; Pizer, S. M.; and Frahm, J.-M. 2019. Recurrent Neural Network for (Un-)Supervised Learning of Monocular Video Visual Odometry and Depth. In *CVPR*.

Wang, T.; Pang, J.; and Lin, D. 2022. Monocular 3D Object Detection with Depth from Motion. In *ECCV*.

Wang, X.; Zhu, Z.; Qin, F.; Ye, Y.; Huang, G.; Chi, X.; He, Y.; and Wang, X. 2022b. MVSTER: Epipolar Transformer for Efficient Multi-View Stereo. In *ECCV*.

Watson, J.; Aodha, O. M.; Prisacariu, V. A.; Brostow, G. J.; and Firman, M. 2021. The Temporal Opportunist: Self-Supervised Multi-Frame Monocular Depth. In *CVPR*.

Wimbauer, F.; Yang, N.; von Stumberg, L.; Zeller, N.; and Cremers, D. 2020. MonoRec: Semi-Supervised Dense Reconstruction in Dynamic Environments from a Single Moving Camera. In *CVPR*.

Xie, J.; Girshick, R.; and Farhadi, A. 2016. Deep3D: Fully Automatic 2D-to-3D Video Conversion with Deep Convolutional Neural Networks. In *ECCV*.

Xu, Q.; and Tao, W. 2019. Learning Inverse Depth Regression for Multi-View Stereo with Correlation Cost Volume. In *AAAI*.

Yang, N.; Wang, R.; Stückler, J.; and Cremers, D. 2018. Deep Virtual Stereo Odometry: Leveraging Deep Depth Prediction for Monocular Direct Sparse Odometry. In *ECCV*.

Yao, Y.; Luo, Z.; Li, S.; Fang, T.; and Quan, L. 2018. MVS-Net: Depth Inference for Unstructured Multi-view Stereo. In *ECCV*.

Yao, Y.; Luo, Z.; Li, S.; Shen, T.; Fang, T.; and Quan, L. 2019. Recurrent MVSNet for High-Resolution Multi-View Stereo Depth Inference. In *CVPR*.

Yin, X.; Wang, X.; Du, X.; and Chen, Q. 2017. Scale Recovery for Monocular Visual Odometry Using Depth Estimated with Deep Convolutional Neural Fields. In *ICCV*.

Zhang, H.; Shen, C.; Li, Y.; Cao, Y.; Liu, Y.; and Yan, Y. 2019. Exploiting temporal consistency for real-time video depth estimation. In *ICCV*.

Zhang, J.; Yao, Y.; Li, S.; Luo, Z.; and Fang, T. 2020. Visibility-aware Multi-view Stereo Network. In *BMVC*.

Zhou, T.; Brown, M.; Snavely, N.; and Lowe, D. G. 2017. Unsupervised Learning of Depth and Ego-Motion from Video. In *CVPR*.

Zhou, Z.; Fan, X.; Shi, P.; and Xin, Y. 2021. R-MSFM: Recurrent Multi-Scale Feature Modulation for Monocular Depth Estimating. In *ICCV*.