

MAMo: Leveraging Memory and Attention for Monocular Video Depth Estimation

Rajeev Yasarla, Hong Cai, Jisoo Jeong, Yunxiao Shi, Risheek Garrepalli, and Fatih Porikli
Qualcomm AI Research*

{ryasarla, hongcai, jisojeon, yunxshi, rgarrepa, fporikli}@qti.qualcomm.com

Abstract

We propose MAMo, a novel memory and attention framework for monocular video depth estimation. MAMo can augment and improve any single-image depth estimation networks into video depth estimation models, enabling them to take advantage of the temporal information to predict more accurate depth. In MAMo, we augment model with memory which aids the depth prediction as the model streams through the video. Specifically, the memory stores learned visual and displacement tokens of the previous time instances. This allows the depth network to cross-reference relevant features from the past when predicting depth on the current frame. We introduce a novel scheme to continuously update the memory, optimizing it to keep tokens that correspond with both the past and the present visual information. We adopt attention-based approach to process memory features where we first learn the spatio-temporal relation among the resultant visual and displacement memory tokens using self-attention module. Further, the output features of self-attention are aggregated with the current visual features through cross-attention. The cross-attended features are finally given to a decoder to predict depth on the current frame. Through extensive experiments on several benchmarks, including KITTI, NYU-Depth V2, and DDAD, we show that MAMo consistently improves monocular depth estimation networks and sets new state-of-the-art (SOTA) accuracy. Notably, our MAMo video depth estimation provides higher accuracy with lower latency, when comparing to SOTA cost-volume-based video depth models.

1. Introduction

Depth plays a fundamental role in 3D perception. Therefore, accurate depth estimation is critical in various applications, such as autonomous driving, AR/VR, and robotics. While it is possible to measure depth using LiDAR or Time-of-Flight (ToF) sensors, these sensors are expensive, consume a lot of power, require extensive calibration, and cannot generate reliable measurements for certain surfaces. On

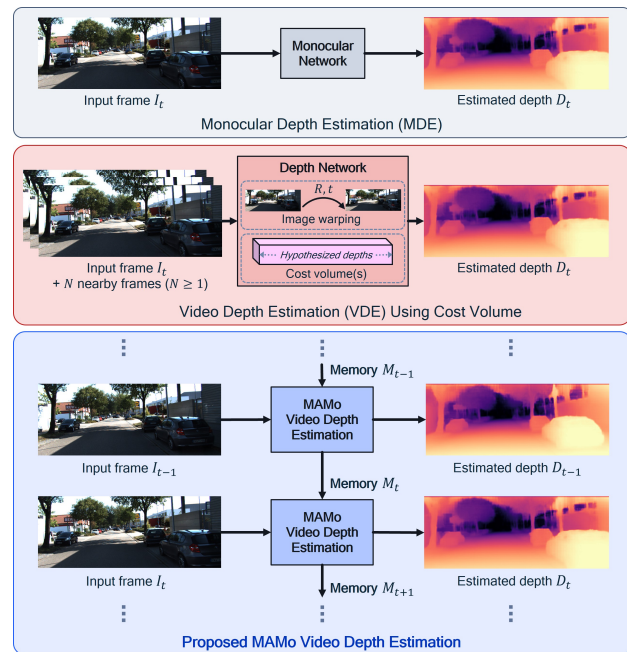


Figure 1. Our proposed MAMo (bottom) enables video depth estimation efficiently in a streaming fashion, by leveraging memory and attention. Monocular depth estimation fails to leverage temporal information (top), while existing cost-volume-based video depth models are computationally expensive (middle). For instance, for each inference, they require multiple image warping operations as well as significant memory usage and heavy computation to construct the cost volume(s).

the other hand, inferring depth from camera images has recently become a cost-efficient and promising alternative. Traditional approaches [42, 14, 35] utilize stereo vision and/or structure-from-motion to estimate depth, which, however, have limited accuracy. By leveraging deep learning, researchers have achieved significantly more accurate image-based depth estimation [11, 13, 3, 38, 62].

Using deep neural networks to infer depth from a single camera image, i.e., monocular depth estimation,¹ has

¹In this paper, we refer to depth estimation based on a single image as *monocular depth estimation* and depth estimation using consecutive frames captured by the same monocular camera as *video depth estimation*.

*Qualcomm AI Research, an initiative of Qualcomm Technologies, Inc.

been one of the most popular choices. Monocular depth estimation, however, only predicts depth based on individual images and does not utilize the temporal information from videos, which are almost always available in many applications, e.g., autonomous driving, AR/VR. More recently, researchers have proposed various ways to leverage multiple frames for depth estimation. One common approach is to utilize a cost volume (or multiple cost volumes), which is used to evaluate depth hypotheses and can be embedded into a deep learning architecture. Cost volumes have enabled considerable boost in performance at the expense of high computational complexity and memory usage. Other works propose video depth estimation models without cost volumes, by leveraging recurrent network [63, 36], optical flow [12, 60], and/or attention [8, 55]. While these models can be more computationally efficient as compared to cost volumes, they have not been shown to provide SOTA accuracy. Moreover, existing video depth estimation methods do not incorporate the latest developments from monocular depth architectures and as a result, they can underperform SOTA monocular depth estimation models despite using more information.

In this paper, we propose a novel approach, MAMo, for video depth estimation, which leverages memory and attention to make use of the key temporal information contained in a video. MAMo can be combined with any monocular network (e.g., NeWCRFs [62], PixelFormer [2]) to perform video depth estimation in a streaming fashion. As such, it is complementary to any existing and future developments in monocular depth estimation. Furthermore, it improves depth estimation accuracy being significantly compute efficient compared to cost volumes.

Fig. 1 (bottom) provides a high-level outline of our proposed MAMo framework. We introduce a memory to augment the depth estimation process as the network goes through the video frames, which maintains learned visual and displacement tokens storing useful information from a set of consecutive previous frames. These tokens are cross-referenced using a cross-attention approach when the network derives the depth for the current input frame.

We propose a novel update scheme for memory module to effectively retain the relevant information from past frames. More specifically, when performing a memory update, we first predict depths using the current frame and a synthesized version of it warped from the previous frame using optical flow, respectively. We compare and minimize the difference between the two predictions, and back-propagate the gradients to update the memory, with the depth network’s weights frozen. Since the memory tokens are used to cross-attend the respective visual features during the two forward passes, they are updated to capture features that is shared across the current frame and the warped previous frame, i.e., the equivariant (*w.r.t.* motion) features

across the current and the previous frames. As we will show, our proposed memory update is more effective compared to sliding window style concatenating

Our main contributions are summarized as follows:

- We introduce MAMo, a novel memory and attention based framework for video depth estimation. MAMo can be combined with any monocular depth network, enabling it to utilize the temporal information to predict more accurate depth.
- In MAMo, we augment model with memory to retain tokens that capture useful information from the previous frames. These tokens are used to assist depth prediction of the current input frame, via cross-attention.
- We propose a novel memory update scheme to effectively retain the relevant information from past frames. Specifically, the memory tokens are updated to encode (motion) equivariant features across the current frame and the previous frame.
- We additionally incorporate careful designs to further improve the video depth estimation performance, such as carrying over decoder features from the previous time step.
- We conduct extensive experiments on common depth estimation datasets: KITTI [17], NYU Depth V2 [46], and DDAD [20]. We show that MAMo not only consistently improves latest monocular depth networks, but also outperforms existing SOTA video depth estimation methods. It is also significantly more efficient as compared to approaches that use cost volumes.

2. Related Work

Monocular Depth Estimation (MDE): In earlier works, various traditional methods have been proposed for monocular depth estimation [33, 34, 41, 43, 54]. Recently, deep-learning-based techniques have gained prominence, which can be broadly categorized into two groups, (i) regressing continuous depth values [18, 62, 64, 6, 45] and (ii) treating depth prediction as classification or ordinal regression [3, 26, 13]. While researchers continue to investigate monocular depth estimation and improve the accuracy, these methods are fundamentally limited as they cannot leverage temporal information when video data is available.

Video Depth Estimation (VDE): Some existing methods devise networks that predict depth based on more than one frame of the video. For instance, ManyDepth [56] utilizes two consecutive frames by leveraging a cost volume. It is also possible to use more frames via cost volume to perform depth estimation, e.g., [30, 44]. However, using more frames can result in delays for the depth prediction, and the cost volume architecture is expensive in terms of computational complexity and memory usage. Other works explore the use of recurrent neural networks, but

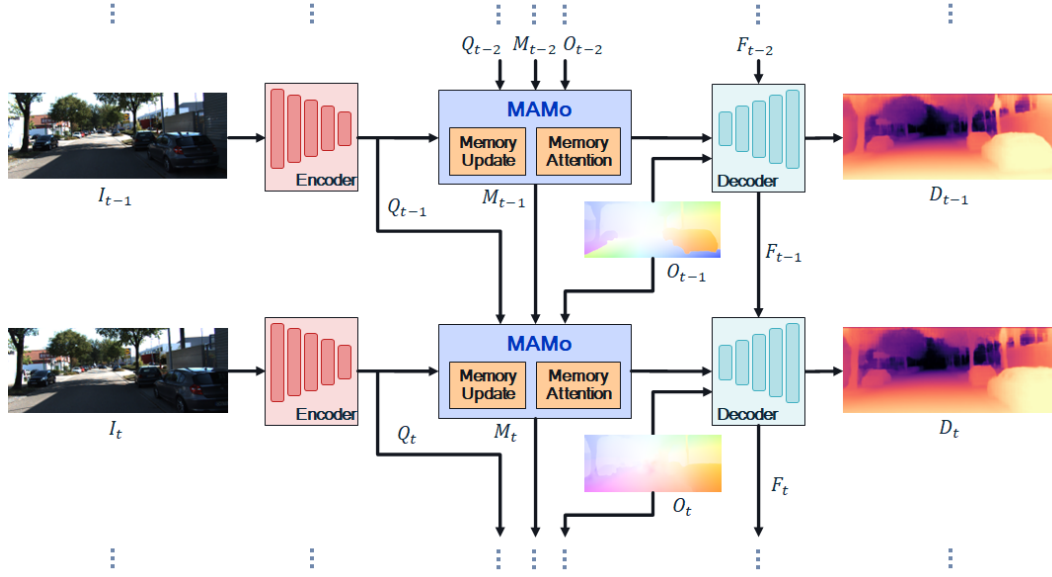


Figure 2. Overview of proposed MAMo method.

only obtain sub-optimal accuracy [12, 63, 36]. More recently, researchers have started to look into leveraging attention mechanisms for video depth estimation, but the existing methods do not achieve state-of-the-art performance, even when compared to the latest MDE models [8, 55]. In the context of video depth estimation, it can be useful to utilize optical flow [50, 21, 48, 16, 23, 22] information to capture the motion across frames. This has been explored by earlier works [12, 60]. Additionally, optical flow driven Depth Estimation with use of memory attention is beneficial for robustness, by using motion cues to detect novel objects i.e., openset and OoD [39] complementing appearance based features, representational learning [28, 15, 29, 10, 5, 4, 52, 51].

Memory: Use of memory techniques is an extensively researched topic in the NLP community [19, 47, 57], addressing reasoning tasks like dialogue communication [58], question-answering [24], and story generation [37]. Memory is recently introduced in computer vision tasks like image captioning [9], colorization [61], text-to-image synthesis [65], video object segmentation [32], and object tracking [7]. Inspired from these methods, we propose MAMo which constructs a memory to maintain relevant spatiotemporal information that can be used to guide depth prediction.

3. Proposed Approach: MAMo

In this section, we present MAMo, a memory and attention based framework for video depth estimation (VDE). We provide an overview of MAMo in Section 3.1. In Sections 3.2 and 3.3, we provide detailed descriptions of the key components of MAMo, i.e., Memory Update (MU) and Memory Attention (MA). In Section 3.4, we discuss additional designs that can further enhance video depth estimation.

We describe training details in Section 3.5.

3.1. Using Memory and Attention for VDE

Consider a sequence of video frames $\{I_0, \dots, I_t, \dots, I_T\}$. We denote the predicted depths on these frames as $\{D_0, \dots, D_t, \dots, D_T\}$ and the estimated optical flows between consecutive frame pairs as $\{O_1, \dots, O_t, \dots, O_T\}$, where O_t is the forward flow from I_{t-1} to I_t .² Given an encoder-decoder architecture for depth estimation, we denote the features extracted by the encoder at time t as Q_t . We denote the encoder as h and the full depth estimation model as g .

Our goal is to develop a depth estimation model that can leverage the temporal correlation across frames as it streams through the video. With this motivation, we augment model with memory to retain a set of learned informative tokens derived from previous L frames as well as the optical flows of the previous L time steps. Formally, $M_t = \{M_t^V, M_t^D\}$ is the memory at time t , and $\{M_t^V\} = \{V_{t-L+1}, \dots, V_t\}$ and $M_t^D = \{P_{t-L+1}, \dots, P_t\}$, where M_t^V stores visual information tokens and M_t^D stores pairwise relative displacement tokens based on the previous L time steps.

At every time step t , given the current input frame I_t , the optical flow O_t , and the previous memory M_{t-1} , MAMo first updates the memory to M_t , in order to capture equivariant information across the current and previous time steps. Next, the updated memory tokens goes through self-attention. The processed features are then fused with the encoder features, Q_t , through cross-attention. The final aggregated features are fed to the decoder to derive the estimated depth, D_t , for the current input frame. Additionally,

²The optical flows can be estimated by models such as RAFT [50].

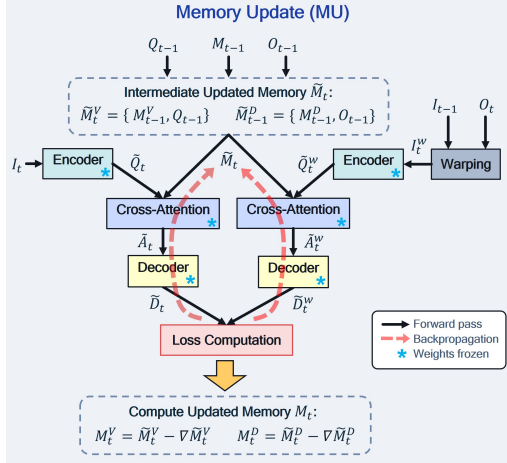


Figure 3. Overview of proposed memory update scheme. To concisely illustrate the main idea of memory update, we omit some operations in the figure, e.g., self-attention on memory tokens (c.f. Section 3.3), decoder feature carry-over (c.f. Section 3.4).

the decoder features from the previous time step, F_{t-1} , are carried over to aid the current prediction. We mathematically define the depth prediction as follows:

$$D_t = g(I_t; M_t, O_t, F_{t-1}). \quad (1)$$

Fig. 2 summarizes our proposed MAMo framework. Overall, it leverages memory and attention mechanisms for video depth estimation. It can be seen that MAMo can readily be implemented on top of any existing monocular depth estimation networks with encoder-decoder architectures. As such, MAMo allows one to take advantage of the latest state-of-the-art and future monocular networks, and convert them to video depth models.

3.2. Memory Update

As the depth estimation model goes through the video frames, it is critical to appropriately update the memory, in order to maintain information that is useful for the current-time depth prediction. As such, we propose a novel scheme to update the memory tokens to capture features that are shared across the current time and the previous time, *i.e.*, equivariant features *w.r.t.* motion across the two frames.

More specifically, at time t , given the previous encoder features Q_{t-1} , memory M_{t-1} , and optical flow O_{t-1} , we first perform an intermediate update to the memory tokens, *i.e.*, $\tilde{M}_t = \{\tilde{M}_t^V, \tilde{M}_t^D\}$, where $\tilde{M}_t^V = \{M_{t-1}^V, Q_{t-1}\}$ which is the concatenation of previous visual tokens and previous encoder features, and $\tilde{M}_t^D = \{M_{t-1}^D, O_{t-1}\}$, which is the concatenation of the previous displacement tokens and previous optical flow. After the concatenation, we discard the first tokens in \tilde{M}_t^V and \tilde{M}_t^D , respectively, so as to maintain the same memory length L .

Next, we perform two forward passes of the depth network, with the network parameters frozen and both using

Algorithm 1 Video depth prediction using MAMo

Input: Video frames $\{I_0, \dots, I_T\}$
Model: $h(\cdot)$ and $g(\cdot)$: encoder and full depth network
Initialization
 $Q_0 \leftarrow h(I_0)$, $O_0 \leftarrow \mathbf{0}$, $F_{-1} \leftarrow \mathbf{0}$
 Update M_0 (repeat Q_0 and O_0 for L times)
 $D_0 \leftarrow g(I_0; M_0, O_0, F_{-1})$
for $I_t \in \{I_1, \dots, I_T\}$ **do**
 Estimate O_t
 Memory Update (Sec. 3.2)
 $\tilde{M}_t^V \leftarrow \{M_{t-1}^V, Q_{t-1}\}$, $\tilde{M}_t^D \leftarrow \{M_{t-1}^D, O_{t-1}\}$
 $\tilde{M}_t \leftarrow \{\tilde{M}_t^V, \tilde{M}_t^D\}$
 $I_t^w \leftarrow \text{Warp}(I_{t-1}, O_t)$
 $\tilde{D}_t \leftarrow g(I_t; \tilde{M}_t, O_t, F_{t-1})$
 $\tilde{D}_t^w \leftarrow g(I_t^w; \tilde{M}_t, O_t, F_{t-1})$
 $\text{SILogLoss}(\tilde{D}_t, \tilde{D}_t^w)$
 Backpropagation
 Update M_t (Eq. 2)
 Depth Estimation
 $D_t \leftarrow g(I_t; M_t, O_t, F_{t-1})$, $Q_t \leftarrow h(I_t)$
end for

the intermediate updated memory. In the first pass, we use the current frame I_t as input and the network predicts a depth map \tilde{D}_t . In the second pass, we construct an input frame, I_t^w , which is warped from I_{t-1} using O_t . In other words, I_t^w is the synthesized version of I_t with motion compensation from O_t . The network consumes I_t^w and generates a depth map \tilde{D}_t^w . We then compute a loss between \tilde{D}_t and \tilde{D}_t^w using the Scale-Invariant Logarithmic (SILog) Loss [11], and backpropagate the gradients to update the memory tokens:

$$M_t^V = \tilde{M}_t^V - \nabla \tilde{M}_t^V, \quad M_t^D = \tilde{M}_t^D - \nabla \tilde{M}_t^D, \quad (2)$$

where $M_t = \{M_t^V, M_t^D\}$ is the updated memory that can then be used to predict the depth map D_t for I_t .

Fig. 3 provides a visual illustration of the memory update steps. During the two forward passes, the memory tokens cross-attend the encoder features \tilde{Q}_t extracted from I_t and the encoder features \tilde{Q}_t^w extracted from I_t^w , respectively. The resulting cross-attended features, \tilde{A}_t and \tilde{A}_t^w , are then fed into the decoder to generate depth predictions, \tilde{D}_t and \tilde{D}_t^w . When we minimize the difference between two outputs, the memory is encouraged to capture features that are shared across Q_t and Q_t^w . This would make \tilde{A}_t and \tilde{A}_t^w more similar since the memory modulates the encoder features via cross-attention, and as a result, \tilde{D}_t and \tilde{D}_t^w would become more similar.

By performing the update, the memory tokens learn to keep similar features shared by I_t and I_t^w . In other words, these are the equivariant features across I_t and I_{t-1} *w.r.t.* the optical flow. Additionally, this update mechanism potentially suppresses the noisy inconsistencies in memory. Overall, our memory and attention framework allows MAMo to be temporally consistent *w.r.t.* motion equivariant features, while also implicitly learning to bet-

ter filter and aggregate spatio-temporal information for non-equivariant/inconsistent regions towards smoother and consistent VDE and hence performing better than sliding window style motion compensated concatenation. We summarize the inference procedure of our proposed MAMo video depth estimation in Algorithm 1.

3.3. Memory Attention

We adopt attention-mechanisms to process updated memory features M_t for the actual depth prediction at time t . First, we perform self-attention over the visual memory tokens M_t^V and also derive corresponding positional encodings from the displacement memory tokens M_t^D . More specifically, we feed M_t^P , which contains the past pairwise optical flow information, into a convolutional block. Linear weighting within convolutional layer operations learns to approximate aggregate estimate of relative motion between the current time and each previous time step tracked in the memory. By doing this, we do not need to explicitly calculate the optical flow between the current time and each of previous time steps which would be computationally demanding. More formally,

$$A_t^{\text{self}} = \text{SelfAttn}(M_t^V; \text{Conv}(M_t^D)), \quad (3)$$

where A_t^{self} are the output features from the self-attention module, $\text{SelfAttn}(x; y)$ denotes self-attention over x with positional encodings of y , and $\text{Conv}(\cdot)$ denotes convolutional layers. Next, the self-attended memory features modulate the encoder features of the current frame via cross-attention:

$$A_t = \text{CrossAttn}(A_t^{\text{self}}, Q_t), \quad (4)$$

where $\text{CrossAttn}(\cdot, \cdot)$ denotes the cross-attention operation and A_t then goes into the decoder for final depth prediction.

3.4. Additional Improvements

To further enable the depth network to utilize available temporal information, we carry over the previous decoder features F_{t-1} and provide them as part of the input to the decoder at time t . In addition the optical flow, O_t , between the previous frame and the current frame is also supplied to the decoder. In this way, the decoder is aware of the relative pixel-wise motion from $t-1$ to t and can thus learn to properly incorporate the previous features for the current depth prediction. While our proposed MAMo approach does not require these additional designs to work well, they do provide further improvements on depth prediction accuracy, as we will show in the experiments.

3.5. Training

In training, at every time-step, we first perform the inference to compute depth (c.f. Algorithm 1), which is then compared with ground-truth depth using SILog loss [11]:

$$\mathcal{L}_d = \alpha \sqrt{\frac{1}{n} \sum_k (\delta d_k)^2 - \frac{\lambda}{n^2} \left(\sum_k \delta d_k \right)^2}, \quad (5)$$

where $\delta d_k = \log D_t(k) - \log D_t^{gt}(k)$, D_t^{gt} is the ground-truth depth for I_t , k is the pixel location, n is the total number of pixels, and $\alpha = 10$ and $\lambda = 0.85$ following [11].

In order to allow the network to train on more motion situations, we employ a video augmentation strategy via sub-sampling. More specifically, we use subsampled sequences of length T for training, i.e., $\{I_0, I_r, \dots, I_{t \times r}, \dots, I_{T \times r}\}$, where r is a sub-sampling ratio randomly selected between 1 and 4 at every epoch for each sequence. As an example, if $r = 4$ and $T = 8$, the video sequence is $\{I_0, I_4, I_8, \dots, I_{32}\}$, this allows the network to see larger motion across the frames. Effectively, this augmentation increases the maximum video frame range to $4 \times T$.

4. Experiments

4.1. Implementation

Networks: We use MAMo to enable latest SOTA monocular methods, e.g., PixelFormer [2], NeWCRFs [62], as well as a strong convolutional baseline, i.e., a variant of DPT [38] with a ResNet encoder (referred to as ResNet-DPT), to perform video depth estimation. Since all these monocular models have encoder-decoder architectures, MAMo can be readily applied. For PixelFormer and NeWCRFs, we extend their own attention designs to create the memory attention modules, respectively. We use Linformer [53] to create the self- and cross-attention blocks for ResNet-DPT.³ To obtain optical flow estimation, we use RAFT [50] for our main results. We also conduct ablation study using the lightweight RAFT-Small model.

Hyperparameters: In all our experiments, we set $T = 8$ and $L = 4$ unless otherwise mentioned. Given the input frame I_t of size $H \times W$, we set the visual memory token size to $512 \times \frac{H}{32} \times \frac{W}{32}$. Thus, the size of M_t^V is $L \times 512 \times \frac{H}{32} \times \frac{W}{32}$ and the size of M_t^D is $L \times 2 \times H \times W$.

Training: We perform all our experiments using 4 NVIDIA-V100 GPUs. We train the network for 25 epochs, using Adam optimizer with a batch size of 8. We set the initial learning rate to 4×10^5 and then linearly decrease it to 4×10^6 across the training iterations.

Evaluation: We use the standard metrics to evaluate depth estimation results; see [11] for metric definitions.

4.2. Datasets

KITTI [17]: KITTI is one of the most commonly used benchmarks for outdoor depth estimation. We use the Eigen split [11] for training and testing, which has 23,488 training images and 697 test images. When training and testing our proposed MAMo approach as well as existing video depth models (e.g., ManyDepth [56]), we use the video (sub)sequences that correspond to the training and test im-

³See supplementary file for more details on how we apply MAMo to these models.

Table 1. Quantitative results on KITTI (Eigen split) for distances up to 80 meters. † means methods uses multiple networks to estimate depth. ManyDepth-FS, and TC-Depth-FS means ManyDepth and TC-Depth are trained in fully-supervised fashion using ground-truths respectively. MF means multi frame methods, SF means single frame methods, and VD means extending MDE to VDE methods. † means higher the better, and ‡ means lower the better.

Type	Method	Encoder	Abs Rel↓	Sq Rel↓	RMSE↓	RMSE _{log} ↓	$\delta < 1.25$ †	$\delta < 1.25^2$ †	$\delta < 1.25^3$ †
MF	NeuralRGB [27]	CNN based†	0.100	–	2.829	–	0.931	–	–
	ST-CLSTM [63]	Resnet18	0.101	–	4.137	–	0.890	0.970	0.9890
	FlowGRU [12]	CNN [12]	0.112	0.700	4.260	0.184	0.881	0.962	0.9830
	Flow2Depth [60]	CNN [31]†	0.081	0.488	3.651	0.146	0.912	0.970	0.9883
	RDE-MV [36]	ResNet18†	0.111	0.821	4.650	0.187	0.821	0.961	0.9823
	Patil <i>et.al.</i> [36]	ResNet18†+ConvLSTM	0.102	–	4.148	–	0.884	0.961	0.9824
	Cao <i>et.al.</i> [8]	–	0.099	–	3.832	–	0.886	0.968	0.9890
	STAD [25]	CNN † [27]	0.109	0.594	3.312	0.153	0.889	0.971	0.9890
	FMNet [55]	ResNeXt-101	0.099	–	3.832	0.129	0.886	0.968	0.9893
	ManyDepth-FS [56]	ResNet50	0.069	0.342	3.414	0.111	0.930	0.989	0.9970
	ManyDepth-FS [56]	Swin-large	0.060	0.248	2.747	0.099	0.955	0.993	0.9981
	TC-Depth-FS [40]	ResNet50	0.071	0.330	3.222	0.108	0.922	0.993	0.9970
SF	AdaBins [3]	EfficientNet-B5+mViT [49]	0.058	0.190	2.360	0.088	0.964	0.995	0.9991
	BinsFormer [26]	Swin-large	0.052	0.151	2.098	0.079	0.975	0.997	0.9992
	DepthFormer [1]	MiT-B4 [59]	0.058	0.187	2.285	0.087	0.967	0.996	0.9991
VD	ResNet-DPT	ResNet50	0.085	0.383	3.242	0.130	0.913	0.981	0.9960
	ResNet-DPT+MAMo (ours)	ResNet50	0.071	0.301	2.984	0.121	0.926	0.990	0.9971
	NeWCRFs [62]	Swin-Base	0.054	0.157	2.140	0.081	0.973	0.997	0.9993
	NeWCRFs+MAMo (ours)	Swin-Base	0.051	0.149	2.090	0.078	0.976	0.998	0.9994
	NeWCRFs	Swin-large	0.053	0.154	2.118	0.080	0.974	0.997	0.9994
	NeWCRFs+MAMo (ours)	Swin-large	0.050	0.141	2.003	0.076	0.977	0.998	0.9994
	PixelFormer [2]	Swin-large	0.052	0.152	2.093	0.079	0.975	0.997	0.9994
	PixelFormer+MAMo (ours)	Swin-large	0.049	0.130	1.884	0.072	0.977	0.998	0.9995

Table 2. Quantitative results on DDAD dataset for distances up to 200 meters, and input frame resolution is 1216 × 1936.

Method	Encoder	Sq Rel↓	RMSE↓	$\delta < 1.25$ †
ManyDepth-FS [56]	ResNet50	5.471	16.123	0.744
ManyDepth-FS [56]	Swin-large	4.211	13.899	0.784
TC-Depth-FS [40]	ResNet50	5.285	15.121	0.777
AdaBins [3]	[49]	4.950	15.228	0.780
DepthFormer [1]	MiT-B4 [59]	4.791	14.595	0.789
ResNet-DPT	ResNet50	5.874	15.697	0.761
ResNet-DPT+MAMo (ours)	ResNet50	4.659	14.288	0.804
NeWCRFs	Swin-large	4.041	11.956	0.816
NeWCRFs+MAMo (ours)	Swin-large	2.990	10.462	0.867
PixelFormer	Swin-large	4.474	12.467	0.802
PixelFormer+MAMo (ours)	Swin-large	3.349	11.094	0.870

ages. The video frames are 375×1241 and the depth range is 80 meters.

DDAD [20]: Dense Depth for Autonomous Driving (DDAD) is very recent dataset featuring urban driving scenarios and long ranges (up to 250 meters). It contains 12,650 training and 3,950 validation samples. We conduct zero-shot transfer to evaluate the generalizability of the trained models from KITTI on all 3,950 validation samples.

NYU Depth V2 [46]: This is a standard dataset for indoor depth estimation, containing 120K RGB-D videos captured from 464 indoor scenes. Since the original test set only contains individual images, we create training and test splits for the video setting. Specifically, from the original 249 training scenes proposed in [11], we use 198 scenes (25,342 image-depth pairs) for training and 86 scenes (10,911 test images) for testing. We refer this video depth version as *NYUDv2-Video*. The images are of size 480×640 with a maximum depth range of 10 meters.

4.3. Results on KITTI

Table 1 shows the results on KITTI. We see that MAMo considerably improves upon ResNet-DPT, NewCRFs, and PixelFormer consistently. We additionally compare our MAMo-based models with existing SOTA monocular depth estimation methods, as well as multi-frame or video depth estimation methods. Note that we retrain the SOTA cost-volume-based and attention-based multi-frame models of ManyDepth [56] and TC-Depth [40] in the supervised setting for a fair comparison; we refer to the supervised versions as ManyDepth-FS and TC-Depth-FS.⁴ It can be seen that our MAMo-based models achieve the SOTA performance and in particular, using MAMo with PixelFormer achieves the best accuracy on KITTI Eigen test set.⁵

4.4. Results on DDAD

In Table 2, we test the KITTI-trained models on DDAD to evaluate the generalization performance. It can be seen that in this case, MAMo still consistently improves the base monocular depth networks. For instance, the squared relative error reduces significantly from 4.041 to 2.990 when applying MAMo to NeWCRFs. When comparing to the other SOTA monocular models as well as the SOTA multi-frame models of ManyDepth-FS and TC-Depth-FS, our MAMo models have the best depth estimation accuracy. This shows that MAMo framework enables the networks to properly utilize the temporal information in the video, allowing them to provide superior generalization ability.

⁴We use the provided supervised learning settings in the original repos.

⁵A more comprehensive comparison table including latest, not officially published methods can be found in the supplementary file.

Table 3. Quantitative results on NYUv2-Video dataset.

Method	Encoder	Abs Rel↓	Sq Rel↓	RMSE↓	RMSE _{log} ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
ResNet-DPT	ResNet50	0.127	0.083	0.464	0.157	0.846	0.974	0.9946
ResNet-DPT + MAMo (ours)	ResNet50	0.119	0.075	0.420	0.135	0.861	0.982	0.9963
NeWCRFs	Swin-Base	0.103	0.056	0.384	0.129	0.900	0.987	0.9978
NeWCRFs + MAMo (ours)	Swin-Base	0.100	0.052	0.371	0.126	0.906	0.989	0.9983
NeWCRFs	Swin-large	0.097	0.051	0.365	0.123	0.912	0.989	0.9981
NeWCRFs + MAMo (ours)	Swin-large	0.094	0.048	0.353	0.120	0.918	0.991	0.9983
PixelFormer	Swin-large	0.098	0.053	0.369	0.124	0.911	0.988	0.9983
PixelFormer + MAMo (ours)	Swin-large	0.094	0.049	0.358	0.119	0.919	0.991	0.9984

Table 4. Ablation comparisons on NYUv2-Video dataset. Here, MA refers to memory attention, MU refers to memory update, OF_d indicates that optical flow is used as input to decoder, and F_{t-1} denotes previous decoder features. OF_m indicates that optical is used to construct memory M_t . Here we use Swin-Base as the encoder for all variants.

Method	F_{i-1}	OF_d	OF_m	MA	MU	Abs Rel↓	Sq Rel↓	RMSE↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$
NeWCRFs						0.1029	0.0559	0.3838	0.9004	0.9872
NeWCRFs + F_{t-1}	✓					0.1014	0.0539	0.3782	0.9021	0.9882
NeWCRFs + warped F_{i-1} using O_t	✓	✓				0.1018	0.0551	0.3779	0.9032	0.9875
NeWCRFs + O_t and F_{t-1} to decoder	✓	✓				0.1009	0.0535	0.3743	0.9049	0.9878
NeWCRFs + sliding window			✓	✓		0.1024	0.0555	0.3827	0.9002	0.9866
NeWCRFs + MU			✓	✓	✓	0.1004	0.0527	0.3722	0.9051	0.9888
NeWCRFs + MAMo (ours)	✓	✓	✓	✓	✓	0.1000	0.0524	0.3708	0.9053	0.9889

Table 5. Using different optical flow networks for MAMo, on NYUv2-Video dataset. We perform this experiment using NeWCRFs + MAMo with Swin-Large encoder.

OF Method	Abs Rel↓	Sq Rel↓	RMSE↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$
RAFT-small	0.095	0.048	0.355	0.917	0.990
RAFT	0.094	0.048	0.353	0.918	0.991

4.5. Results on NYUv2-Video

To assess the benefits of using our proposed MAMo in indoor scenarios, we train and evaluate the performance of PixelFormer [2], NeWCRFs [62], and ResNet-DPT [38], with and without MAMo on the NYUv2-Video dataset. Table 3 shows that in the indoor setting, our proposed MAMo approach is useful for improve the accuracy of monocular models by properly leveraging the video information.

5. Discussion

5.1. Ablation Study

We conduct extensive ablation studies to analyze different aspects of our proposed approach and the design choices. We conduct experiments on NYUv2-Video using NeWCRFs as the base method, with the Swin-Base encoder. As shown in Table 4, we compare the following options: (1) original NeWCRFs, (2) NeWCRFs + F_{t-1} : previous decoder features are used, (3) NeWCRFs + warp F_{i-1} using O_t : we warp F_{t-1} using O_t before passing it to the decoder, (4) NeWCRFs + O_t and F_{t-1} to the decoder: O_t and F_{t-1} are concatenated along with encoder features as input to the decoder, (5) NeWCRFs + sliding window: we construct M_t using sliding-window technique to save previous features and optical flows, and then use attention to fuse them with Q_t , (6) NeWCRFs + MU: we update M_t using our proposed memory update scheme (c.f. Section 3.2), before feeding it to memory attention and decoder, (7) NeWCRFs + MAMo:

Table 6. Using different memory length L for MAMo. We perform this study using NeWCRFs + MAMo with Swin-Large encoder.

dataset	L	Abs Rel↓	Sq Rel↓	RMSE↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$
KITTI	2	0.051	0.149	2.030	0.976	0.998
	4	0.050	0.141	2.003	0.977	0.998
	6	0.050	0.140	1.990	0.977	0.999
NYUv2 (video)	2	0.095	0.049	0.358	0.917	0.990
	4	0.094	0.048	0.353	0.918	0.991
	6	0.093	0.048	0.351	0.920	0.991

applying our full proposed MAMo to NewCRFs.

It can be seen in Table 4, carrying over F_{t-1} provides a 1.5% decrease in RMSE. On the other hand, warping F_{t-1} using O_t is not helpful and incurs additional computation to warp the features at multiple scales. Using F_{t-1} and O_t as input to decoder via concatenation, however, enables the decoder to learn efficient positional or motion cues for F_{t-1} w.r.t. I_t , thus decreasing RMSE by 2.5%. Constructing the memory M_t naively with sliding window does not improve the depth estimation accuracy, since this approach keeps irrelevant features in M_t , making it more difficult for attention to fuse M_t with Q_t ; sliding window results in lower accuracy as shown in Table 4. In contrast, our proposed memory update approach (c.f. Section 3.2) allows the memory to maintain only highly correlated information and reduces RMSE by 3.5% and squared relative error by 6%. Finally, combining all our proposed techniques, our full MAMo approach considerably decreases the squared relative error by 6.5% and RMSE by 4% when comparing to the NeWCRFs.

Optical Flow: While we use RAFT to estimate optical flow in our main experiments, MAMo still works well with lighter, more efficient optical flow networks. To show this, we compare the depth estimation performance of MAMo by using optical flows generated by RAFT and RAFT-Small. RAFT-Small has significantly lower latency; see [50] for




Figure 4. Qualitative results on KITTI. We highlight (white boxes) regions where MAMo significantly improves depth estimation quality.

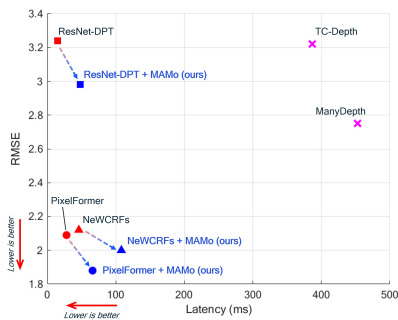


Figure 5. RMSE vs. Latency on KITTI Eigen test set.

detailed comparison. It can be seen in Table 5, when using lighter-weight optical flow models like RAFT-small, the depth estimation performance is almost the same as that of using RAFT.

Memory length L : We perform experiment to ablate on different memory lengths L (e.g., 2, 4, and 6) on KITTI and NYUv2-Video. Note that we set $T = 8$ in these experiments. In Table 6, we can see the depth estimation accuracy considerably improves from $L = 2$ to $L = 4$. Increasing it further to $L = 6$ provides additional minor improvements.

5.2. Computation: Accuracy vs. Efficiency

We compare the average inference time using our proposed MAMo models with existing state-of-the-art multi-frame video depth estimation methods on KITTI. In Fig. 5, we see that MAMo enables considerable accuracy improvements while incurring minor additional runtime. On the other hand, ManyDepth and TC-Depth have significantly larger latencies. For instance, ManyDepth has a latency of over 450 ms while PixelFormer + MAMo requires 66 ms, and PixelFormer + MAMo provides significantly better accuracy as compared to ManyDepth.

Although we do not include the optical flow inference times in the figure, the latency of modern optical flow models are not large. For instance, RAFT and RAFT-small have latencies of 115 ms and 60 ms on KITTI images, respectively. As such, even in the case where the optical flow and depth models are run in a sequential manner, the MAMo-based models still have significantly lower latencies as compared to ManyDepth and TC-Depth. Latencies are measured using a 11GB NVIDIA RTX-2080 GPU.

5.3. Qualitative Results on KITTI

Fig. 4 shows that MAMo considerably improves depth estimation over baselines PixelFormer and NeWCRFs in several regions, e.g., the traffic sign and the building facade in the left sample, the biker in the middle sample, and the fences in the right sample (highlighted by white boxes). Overall, MAMo provides clearer and sharper depth maps.

6. Conclusions

In this paper, we proposed a novel monocular video depth estimation approach, MAMo, which leverages memory and attention, and can be applied to any existing monocular depth estimation networks to transform them into video prediction models. Specifically, in MAMo, we propose a novel memory update method that allows the memory to maintain relevant and useful information for depth estimation, as the model goes through a video. We devise attention schemes to combine the information from the memory with the visual features of the current time, in order to predict accurate depth for the current input frame. Our extensive results and ablation study on: KITTI, NYU Depth V2, and DDAD, confirms that our MAMo approach is effective, improving monocular depth estimation accuracy consistently.

References

- [1] Ashutosh Agarwal and Chetan Arora. Depthformer: Multi-scale vision transformer for monocular depth estimation with global local information fusion. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 3873–3877, 2022. 6
- [2] Ashutosh Agarwal and Chetan Arora. Attention attention everywhere: Monocular depth prediction with skip attention. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5861–5870, January 2023. 2, 5, 6, 7
- [3] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4009–4018, 2021. 1, 2, 6
- [4] Shubhankar Borse, Debasmit Das, Hyojin Park, Hong Cai, Risheek Garrepalli, and Fatih Porikli. Dejavu: Conditional regenerative learning to enhance dense prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19466–19477, June 2023. 3
- [5] Shubhankar Borse, Hyojin Park, Hong Cai, Debasmit Das, Risheek Garrepalli, and Fatih Porikli. Panoptic, instance and semantic relations: A relational context encoder to enhance panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1269–1279, June 2022. 3
- [6] Hong Cai, Janarbek Matai, Shubhankar Borse, Yizhe Zhang, Amin Ansari, and Fatih Porikli. X-distill: Improving self-supervised monocular depth via cross-task distillation. In *British Machine Vision Conference*, 2021. 2
- [7] Jiarui Cai, Mingze Xu, Wei Li, Yuanjun Xiong, Wei Xia, Zhuowen Tu, and Stefano Soatto. Memot: multi-object tracking with memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8090–8100, 2022. 3
- [8] Yuanzhouhan Cao, Yidong Li, Haokui Zhang, Chao Ren, and Yifan Liu. Learning structure affinity for video depth estimation. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 190–198, 2021. 2, 3, 6
- [9] Cesc Chunseong Park, Byeongchang Kim, and Gunhee Kim. Attend to you: Personalized image captioning with context sequence memory networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 895–903, 2017. 3
- [10] Debasmit Das, Shubhankar Borse, Hyojin Park, Kambiz Azarian, Hong Cai, Risheek Garrepalli, and Fatih Porikli. Transadapt: A transformative framework for online test time adaptive semantic segmentation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 3
- [11] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014. 1, 4, 5, 6
- [12] Chanho Eom, Hyunjong Park, and Bumsub Ham. Temporally consistent depth prediction with flow-guided memory units. *IEEE Transactions on Intelligent Transportation Systems*, 21(11):4626–4636, 2019. 2, 3, 6
- [13] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2002–2011, 2018. 1, 2
- [14] Yasutaka Furukawa, Brian Curless, Steven M Seitz, and Richard Szeliski. Towards internet-scale multi-view stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1434–1441, 2010. 1
- [15] Risheek Garrepalli. Oracle analysis of representations for deep open set detection. *arXiv preprint arXiv:2209.11350*, 2022. 3
- [16] Risheek Garrepalli, Jisoo Jeong, Rajeswaran C. Ravindran, Jamie Menjay Lin, and Fatih Porikli. Dift: Dynamic iterative field transforms for memory efficient optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2219–2228, June 2023. 3
- [17] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361, 2012. 2, 5
- [18] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [19] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014. 3
- [20] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3D Packing for Self-Supervised Monocular Depth Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 6
- [21] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer: A transformer architecture for optical flow. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*, pages 668–685. Springer, 2022. 3
- [22] Jisoo Jeong, Hong Cai, Risheek Garrepalli, and Fatih Porikli. Distractflow: Improving optical flow estimation via realistic distractions and pseudo-labeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13691–13700, June 2023. 3
- [23] Jisoo Jeong, Jamie Menjay Lin, Fatih Porikli, and Nojun Kwak. Imposing consistency for optical flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3181–3191, 2022. 3
- [24] Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain

- Paulus, and Richard Socher. Ask me anything: Dynamic memory networks for natural language processing. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1378–1387, 2016. 3
- [25] Hyunmin Lee and Jaesik Park. Stad: Stable video depth estimation. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 3213–3217. IEEE, 2021. 6
- [26] Zhenyu Li, Xuyang Wang, Xianming Liu, and Junjun Jiang. Binsformer: Revisiting adaptive bins for monocular depth estimation. *arXiv preprint arXiv:2204.00987*, 2022. 2, 6
- [27] Chao Liu, Jinwei Gu, Kihwan Kim, Srinivasa G Narasimhan, and Jan Kautz. Neural rgb (r) d sensing: Depth and uncertainty from a video camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10986–10995, 2019. 6
- [28] Si Liu, Risheek Garrepalli, Thomas Dietterich, Alan Fern, and Dan Hendrycks. Open category detection with PAC guarantees. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3169–3178. PMLR, 10–15 Jul 2018. 3
- [29] Si Liu, Risheek Garrepalli, Dan Hendrycks, Alan Fern, Debashis Mondal, and Thomas G Dietterich. Pac guarantees and effective algorithms for detecting novel categories. *J. Mach. Learn. Res.*, 23:44–1, 2022. 3
- [30] Xiaoxiao Long, Lingjie Liu, Wei Li, Christian Theobalt, and Wenping Wang. Multi-view depth estimation using epipolar spatio-temporal networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8258–8267, June 2021. 2
- [31] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4040–4048, 2016. 6
- [32] Jiayu Miao, Yunchao Wei, and Yi Yang. Memory aggregation networks for efficient interactive video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10366–10375, 2020. 3
- [33] Jeff Michels, Ashutosh Saxena, and Andrew Y Ng. High speed obstacle avoidance using monocular vision and reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 593–600, 2005. 2
- [34] Takaaki Nagai, Takumi Naruse, Masaaki Ikehara, and Akira Kurematsu. Hmm-based surface reconstruction from single images. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, volume 2, pages II–II. IEEE, 2002. 2
- [35] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. DTAM: Dense tracking and mapping in real-time. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 2320–2327. IEEE, 2011. 1
- [36] Vaishakh Patil, Wouter Van Gansbeke, Dengxin Dai, and Luc Van Gool. Don’t forget the past: Recurrent depth estimation from monocular video. *IEEE Robotics and Automation Letters*, 5(4):6813–6820, 2020. 2, 3, 6
- [37] Tanzila Rahman, Hsin-Ying Lee, Jian Ren, Sergey Tulyakov, Shweta Mahajan, and Leonid Sigal. Make-a-story: Visual memory conditioned consistent story generation. *arXiv preprint arXiv:2211.13319*, 2022. 3
- [38] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12179–12188, 2021. 1, 5, 7
- [39] Lukas Ruff, Jacob R Kauffmann, Robert A Vandermeulen, Grégoire Montavon, Wojciech Samek, Marius Kloft, Thomas G Dietterich, and Klaus-Robert Müller. A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109(5):756–795, 2021. 3
- [40] Patrick Ruhkamp, Daoyi Gao, Hanzhi Chen, Nassir Navab, and Benjamin Busam. Attention meets geometry: Geometry guided spatial-temporal attention for consistent self-supervised monocular depth estimation. In *Proceedings of the International Conference on 3D Vision (3DV)*, pages 837–847, 2021. 6
- [41] Ashutosh Saxena, Sung Chung, and Andrew Ng. Learning depth from single monocular images. *Advances in Neural Information Processing Systems*, 18, 2005. 2
- [42] Ashutosh Saxena, Jamie Schulte, Andrew Y Ng, et al. Depth estimation using monocular and stereo cues. In *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI)*, volume 7, pages 2197–2203, 2007. 1
- [43] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):824–840, 2008. 2
- [44] Mohamed Sayed, John Gibson, Jamie Watson, Victor Prisacariu, Michael Firman, and Clément Godard. Simplerecon: 3d reconstruction without 3d convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [45] Yunxiao Shi, Hong Cai, Amin Ansari, and Fatih Porikli. Ega-depth: Efficient guided attention for self-supervised multi-camera depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 119–129, 2023. 2
- [46] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 7576, pages 746–760, 2012. 2, 6
- [47] Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. End-to-end memory networks. *Advances in Neural Information Processing Systems*, 28, 2015. 3
- [48] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8934–8943, 2018. 3
- [49] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Proceedings of*

- the International Conference on Machine Learning (ICML)*, pages 6105–6114. PMLR, 2019. [6](#)
- [50] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 402–419. Springer, 2020. [3](#), [5](#), [7](#)
- [51] Vibashan VS, Poojan Oza, and Vishal M Patel. Towards on-line domain adaptive object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 478–488, 2023. [3](#)
- [52] Vibashan VS, Ning Yu, Chen Xing, Can Qin, Mingfei Gao, Juan Carlos Niebles, Vishal M Patel, and Ran Xu. Mask-free ovis: Open-vocabulary instance segmentation without manual mask annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23539–23549, 2023. [3](#)
- [53] Sinong Wang, Belinda Z Li, Madian Khabisa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. [5](#)
- [54] Xiaoyan Wang, Chunping Hou, Liangzhou Pu, and Yonghong Hou. A depth estimating method from a single image using foe crf. *Multimedia Tools and Applications*, 74:9491–9506, 2015. [2](#)
- [55] Yiran Wang, Zhiyu Pan, Xingyi Li, Zhiguo Cao, Ke Xian, and Jianming Zhang. Less is more: Consistent video depth estimation with masked frames modeling. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 6347–6358, 2022. [2](#), [3](#), [6](#)
- [56] Jamie Watson, Oisín Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1164–1174, 2021. [2](#), [5](#), [6](#)
- [57] Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014. [3](#)
- [58] Chien-Sheng Wu, Richard Socher, and Caiming Xiong. Global-to-local memory pointer networks for task-oriented dialogue. *arXiv preprint arXiv:1901.04713*, 2019. [3](#)
- [59] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. [6](#)
- [60] Jiaxin Xie, Chenyang Lei, Zhuwen Li, Li Erran Li, and Qifeng Chen. Video depth estimation by fusing flow-to-depth proposals. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 10100–10107, 2020. [2](#), [3](#), [6](#)
- [61] Seungjoo Yoo, Hyojin Bahng, Sunghyo Chung, Junsoo Lee, Jaehyuk Chang, and Jaegul Choo. Coloring with limited data: Few-shot colorization via memory augmented networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11283–11292, 2019. [3](#)
- [62] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. Newcrfs: Neural window fully-connected crfs for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [1](#), [2](#), [5](#), [6](#), [7](#)
- [63] Haokui Zhang, Chunhua Shen, Ying Li, Yuanzhouhan Cao, Yu Liu, and Youliang Yan. Exploiting temporal consistency for real-time video depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1725–1734, 2019. [2](#), [3](#), [6](#)
- [64] Jing Zhu, Yunxiao Shi, Mengwei Ren, and Yi Fang. Mda-net: memorable domain adaptation network for monocular depth estimation. In *British Machine Vision Conference 2020*, 2020. [2](#)
- [65] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5802–5810, 2019. [3](#)