

Learning Monocular Depth in Dynamic Environment via Context-aware Temporal Attention

Zizhang Wu^{1*}, Zhuozheng Li^{1*}, Zhi-Gang Fan^{1*}, Yunzhe Wu¹, Yuanzhu Gan¹ and Jian Pu²

¹ZongmuTech

²Fudan University

wuzizhang87@gmail.com, {zhuozheng.li, zhigang.fan, nelson.wu, yuanzhu.gan}@zongmotech.com
jianpu@fudan.edu.cn

Abstract

The monocular depth estimation task has recently revealed encouraging prospects, especially for the autonomous driving task. To tackle the ill-posed problem of 3D geometric reasoning from 2D monocular images, multi-frame monocular methods are developed to leverage the perspective correlation information from sequential temporal frames. However, moving objects such as cars and trains usually violate the static scene assumption, leading to feature inconsistency deviation and misaligned cost values, which would mislead the optimization algorithm. In this work, we present **CTA-Depth**, a **Context-aware Temporal Attention** guided network for multi-frame monocular **Depth** estimation. Specifically, we first apply a multi-level attention enhancement module to integrate multi-level image features to obtain an initial depth and pose estimation. Then the proposed CTA-Refiner is adopted to alternatively optimize the depth and pose. During the CTA-Refiner process, context-aware temporal attention (CTA) is developed to capture the global temporal-context correlations to maintain the feature consistency and estimation integrity of moving objects. In particular, we propose a long-range geometry embedding (LGE) module to produce a long-range temporal geometry prior. Our approach achieves significant improvements (e.g., 13.5% for the Abs Rel metric on the KITTI dataset) over state-of-the-art approaches on three benchmark datasets.

1 Introduction

Monocular depth estimation aims at predicting accurate pixel-wise depth from monocular RGB images. Due to its low cost and easy implementation, monocular depth estimation has achieved promising prospects in practical applications [Li *et al.*, 2022; Wu *et al.*, 2022; Mumuni and Mumuni, 2022]. Particularly, monocular depth estimation [Li *et al.*, 2015; Liu *et al.*, 2015; Ricci *et al.*, 2018; Bhat *et al.*, 2021; Yuan *et al.*, 2022] under the single-frame setting has achieved convincing results by conducting robust convolutional neural networks with prior geometric constraints. Nevertheless,

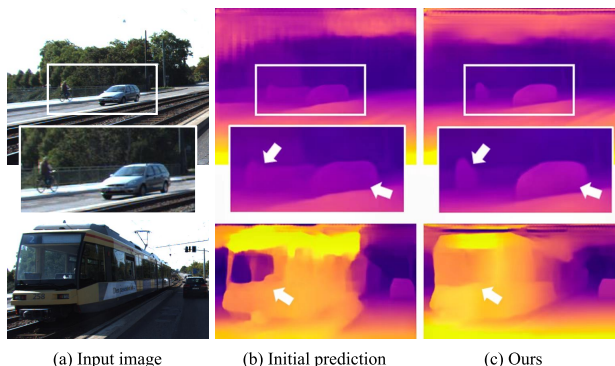


Figure 1: Given input images, our CTA-Depth with CTA-Refiner predicts more accurate depth maps compared to the initial prediction, especially for dynamic objects.

it is still challenging to precisely recover the 3D environment from a single monocular 2D image. On the other hand, noting that sequential image frames are achievable from the monocular camera, existing studies [Wang *et al.*, 2019; Zhang *et al.*, 2019; Patil *et al.*, 2020] start paying greater attention to the depth estimation under the multi-frame setting. Inspired by the stereo matching task, multi-frame monocular depth estimation works typically employed the cost volume or cost map [Watson *et al.*, 2021; Gu *et al.*, 2023; Bae *et al.*, 2022] to accomplish geometric reasoning and have gradually achieved state-of-the-art performance.

However, the widely applied static scene assumption [Klingner *et al.*, 2020; Li *et al.*, 2019] for the construction of the cost volume does not always hold in real-world scenarios. Specifically, moving objects such as cars, trains and pedestrians result in feature inconsistency deviation, misaligned cost values and degraded re-projection loss, which would mislead the optimization algorithm. To address this issue, recent works [Lee *et al.*, 2021a; Feng *et al.*, 2022; Wimbauer *et al.*, 2021] attempted to solve dynamic problems by introducing an auxiliary object motion prediction module and segmentation masks to predict or disentangle dynamic objects explicitly. It inevitably increases the complexity and redundancy of the model and ignores the sustained temporal relation modeling of moving objects across long-range multiple frames, which thus limits the potential of sequential images for time-crucial industry implementation.

Thus, to efficiently boost the multi-frame context-temporal feature integration for dynamic targets without explicit auxiliary modules, we propose our **CTA-Depth**, a Context-aware Temporal Attention network for joint multi-frame monocular **Depth** and pose estimation. Specifically, we first utilize the multi-level attention enhancement (MAE) module for reliable initial estimation, which applies cross-scale attention layers to achieve ample interaction of different-scale features to equip the network with both local and global attentions. Furthermore, we develop the refiner CTA-Refiner to iteratively optimize our predictions with the inputs of context and temporal features. In specific, we develop our depth-pose context-aware temporal attention (CTA) with the cross-attention mechanism that assigns the temporal features as values, and the context features as queries and keys. As a result, it implicitly interacts with the context and temporal features to maintain the estimation integrity of moving objects across multiple sample frames. Additionally, to expand the temporal field of interest and aggregate useful clues for geometry reasoning, especially for dynamic targets within distant frames, we present our long-range geometry embedding (LGE) and provide it to the CTA process for seizing the long-range temporal geometry prior.

Qualitatively, as shown in Fig. 1(b), it is difficult for the single-frame method to recover the complete geometry of moving objects, while the predictions from our implicit CTA-Depth in Fig. 1(c) demonstrate its robustness at high-level feature recognition. Quantitatively, we conduct extensive experiments on three challenging benchmarks to validate the effectiveness of our pipeline against state-of-the-art models. We summarize our contributions as follows:

- We propose our CTA-Depth, an implicit, long-range Context-aware Temporal Attention guided network for supervised multi-frame monocular depth estimation, focusing on the dynamic object areas. It achieves state-of-the-art performance on challenging KITTI, VKITTI2, and nuScenes datasets.
- We design a novel depth-pose context-aware temporal attention (CTA), which implicitly learns the temporal geometric constraints for moving objects via attention-based integration.
- We introduce a novel long-range geometry embedding (LGE) module to promote geometry reasoning among the long-range temporal frames.
- We develop an effective multi-level attention enhancement (MAE) module to make global-aware initial depth and estimations. It promotes the distinction of far-away small objects from the static background.

2 Related Work

2.1 Monocular Depth Estimation

Convolutional neural networks with LiDAR supervision [Wang *et al.*, 2015; Fu *et al.*, 2018; Tang and Tan, 2018; Teed and Deng, 2019; Lee and Kim, 2020; Guizilini *et al.*, 2021; Lee *et al.*, 2021b] have shown promising results in monocular depth estimation. As a pioneer, Eigen *et al.* [Eigen

et al., 2014] directly regressed depth by employing two stacked deep networks that made a coarse prediction from the whole image and then refined it locally. On the other hand, [Laina *et al.*, 2016] adopted an end-to-end single CNN architecture with residual learning. To guide the encoded features to the desired depth prediction, [Lee *et al.*, 2019] further deployed it with local planar guidance layers. Recently, [Ranftl *et al.*, 2021] introduced a dense prediction transformer [Vaswani *et al.*, 2017] for depth prediction. Meanwhile, [Bhat *et al.*, 2021; Yang *et al.*, 2021] developed global information processing with vision transformer [Dosovitskiy *et al.*, 2021] for performance boost. Besides, [Yuan *et al.*, 2022] adopted the swin-transformer [Liu *et al.*, 2021] as the image encoder and the neural window fully-connected conditional random fields (NeWCRFs) module as the feature decoder. In particular, inspired by the RAFT [Teed and Deng, 2020] which employed a GRU-based recurrent operator to update optical flow, [Gu *et al.*, 2023] proposed a multi-frame monocular approach with a deep recurrent optimizer to update the depth and camera poses alternately.

However, these cost-map-based multi-frame methods [Teed and Deng, 2019; Gu *et al.*, 2023] lead to performance degradation within dynamic areas due to the static scene assumption. To solve this problem, we introduce a long-range geometry embedding module and effectively inject the proposed depth-pose context-aware temporal attention into the deep refinement network for the optimization process of depth and pose.

2.2 Depth Estimation in Dynamic Environment

Moving objects significantly hamper the multi-frame matching strategy due to the inevitable object truncation and occlusion. Specifically, both the re-projection loss calculation and the cost volume construction fall into catastrophic failure cases during the positional change of observed targets. Existing works [Li *et al.*, 2019; Lee *et al.*, 2021a; Watson *et al.*, 2021] thus leveraged the segmentation mask to separate the static-scene depth loss from moving objects. In particular, they also proposed explicit object motion prediction and a disentanglement module to assist the cost volume construction. Specifically, SGDepth [Klingner *et al.*, 2020] proposed a semantic masking scheme providing guidance to prevent dynamic objects from contaminating the photometric loss. DynamicDepth [Feng *et al.*, 2022] introduced an object motion disentanglement module that takes dynamic category segmentation masks as input to explicitly disentangle dynamic objects. Considering time-crucial tasks such as autonomous driving, instead of adopting explicit redundant static-dynamic separation algorithms, we developed an efficient implicit modeling pipeline with our novel context-aware temporal attention module. Besides, noting that previous works limited their methods to a few frames only due to an increase in computational cost, we developed a novel geometry embedding module to effectively encode semantic guidance from long-range time series.

As a result, our pipeline can dynamically interact with the long-range semantic flow with the current-frame spatial geometry in a fully differentiable manner and is thus available for real industry implementation.

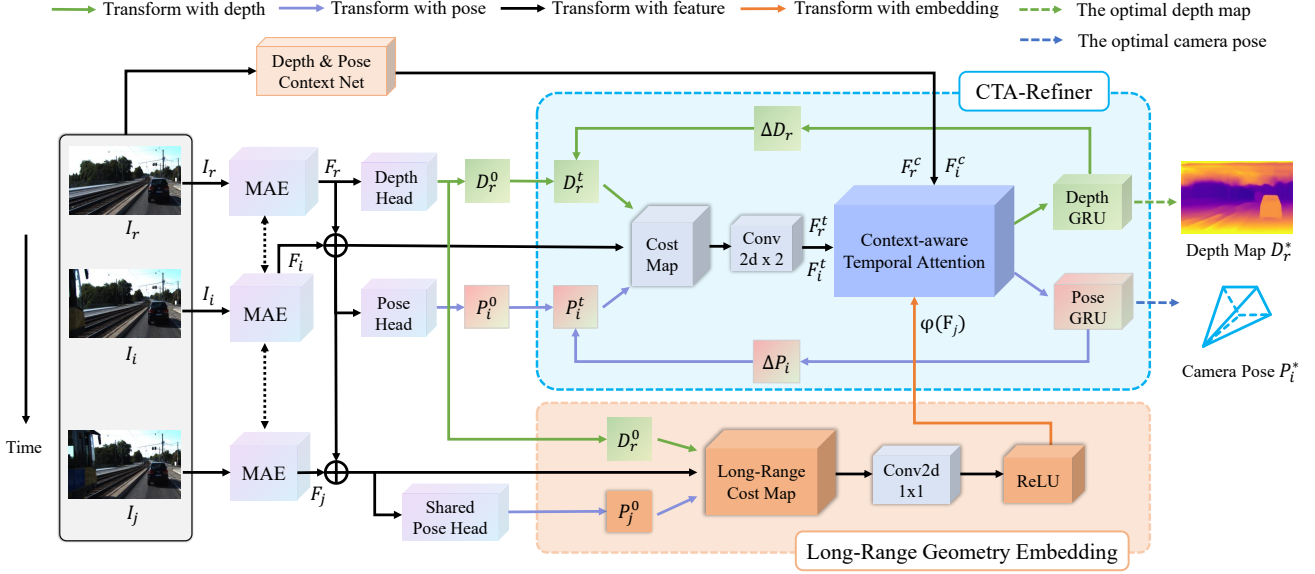


Figure 2: Overview of our CTA-Depth. Given a reference image I_r and its N monocular sequential images $\{I_i\}_{i=1}^N$, we first group N pairs of network inputs, where each pair is composed of I_r and a sequential image I_i . We then feed the two images into the MAE module and go through two heads to obtain the initial depth D_r^0 , and pose P_i^0 . Hence, the CTA-Refiner is proposed to alternately update the depth map and pose through iterations until the optimum solution, D_r^* and P_i^* . In particular, with multiple temporal-neighboring frames I_j , we also design a long-range geometry embedding module to provide long-range temporal geometric priors for the depth refiner efficiently.

3 Method

3.1 Overview

We demonstrate the framework of our approach in Fig. 2, which mainly consists of five components: the network inputs, the multi-level attention enhancement (MAE) module, the depth & pose context net, and the context-aware temporal attention refiner (CTA-Refiner) which includes the depth-pose context-aware temporal attention (CTA) module and the last long-range geometry embedding (LGE) module.

Given a reference image I_r within one video captured via a monocular camera, and its N video frames $\{I_i\}_{i=1}^N$, our goal is to predict the I_r 's accurate depth D_r^* and the relative camera poses $\{P_i^*\}_{i=1}^N$ for sequence image I_i with respect to I_r . Specifically, we first regard the monocular video frames as network inputs. Then, we adopt the multi-level attention enhancement (MAE) module to extract representative visual features for the following depth head and (shared) pose head, which produces the initial depth D_r^0 and initial pose P^0 . Meanwhile, we employ the depth & pose context net to extract context features F_r^c for depth and F_i^c for pose. In addition, the long-range geometry embedding module seizes the multi-frame temporal geometry knowledge to create the long-range geometry embedding (LGE) $\varphi(F_j)$, which is further provided to the CTA to assist the refinement. Afterwards, we adopt the context-aware temporal attention refiner (CTA-Refiner) to alternately update the results, which gradually converge to the optimal depth D_r^* and pose P_i^* .

3.2 Multi-level Attention Enhancement (MAE)

Affiliated with the optimizer-based pipeline, we deliver our multi-level attention enhancement (MAE) module to achieve the initial prediction of depth and pose. As shown in Fig. 3,

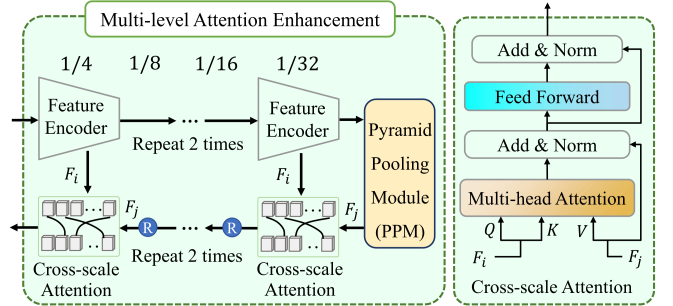


Figure 3: The architecture of our multi-level attention enhancement (MAE) module, which adequately integrates multi-level image features via feature encoders, cross-scale attention layers, and PPM module. “R” denotes the rearranged up-scaling for feature maps.

we propose the multi-level feature setting and pyramid pooling module to reinforce the interest of the far-away small targets. In addition, to distinguish distant moving objects from the static background, we adopt cross-scale attention layers to enhance the interaction of different-scale features.

Specifically, as shown in Fig. 3, we utilize the feature encoder to extract four different-scale features from I_r . The low-level features focus on the local details, while the high-level features seize the global context information, which both contribute to the rising interest of the distant targets [Lin *et al.*, 2017; Liu *et al.*, 2018]. Afterwards, we employ the pyramid pooling module (PPM) [Zhao *et al.*, 2017] to aggregate these features and deliver four cross-scale attention layers to fuse the multi-level feature maps. We use the scaled features F_i as the query and key, and adopt the fusion features F_j as the value to stimulate the interaction of multi-scale features. Within each cross-attention layer, we also introduce the

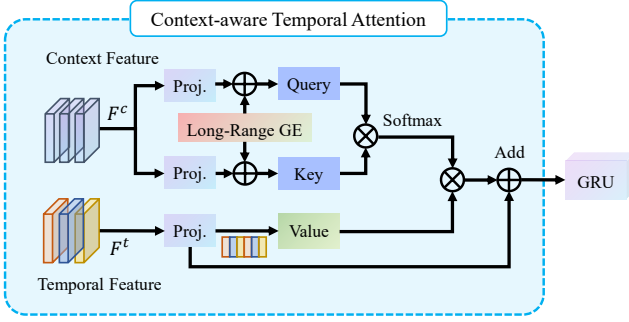


Figure 4: Illustration of the Depth/Pose CTA. Our CTA-Refiner alternately optimizes the depth and the pose. During the depth refinement, the depth CTA uses the depth context feature F_r^c , the temporal feature F_r^t , and the Long-Range GE as inputs and feeds the outputs into the depth GRU. For the pose refinement, the pose CTA employs the pose context-aware feature F_i^c , the temporal feature F_i^t and the Long-Range GE as inputs, then its outputs are fed into pose GRU.

rearranged up-scaling operation [Yuan *et al.*, 2022] to reduce the network complexity and boundary contour refinements. As shown in Fig. 5, our method achieves accurate depth prediction of objects at different scales and long distances.

3.3 Context-aware Temporal Attention (CTA)

CTA-Refiner

We adopt the CTA-Refiner to iteratively refine initial estimations to the final converged depth D_r^* and pose P_i^* , together with two introduced extra inputs: the context features F_r^c for depth and F_i^c for the pose from the depth & pose context net. The refiner accepts these inputs to produce the prediction offset ΔD_r^t and ΔP_i^t and then updates the results as follows:

$$D_r^{t+1} \leftarrow D_r^t + \Delta D_r^t, \quad (1)$$

$$P_i^{t+1} \leftarrow P_i^t + \Delta P_i^t. \quad (2)$$

In detail, we first calculate the cost map given D_r^t , P_i^t , F_r and F_i , as shown in Fig. 2, where optimization for pose remains freezing when optimizing depth, and vice versa. Notably, the cost map measures the distance in feature space between I_r and the sequence image I_i . Next, we adopt a simple feature extractor (two convolutions) to obtain the temporal features F_r^t from the cost map, preparing for the following depth-pose CTA. Thus it implicitly rectifies the implied content-temporal inconsistency for moving objects and effectively promotes information integration between temporal features and depth-pose context features.

We formulate the cost map as the $L2$ distance between aligned feature maps \mathcal{F} . Given the depth map D of the reference image I_r and the relative camera pose T_i of another image I_i with respect to I_r , the cost is constructed at each pixel x in the reference image I_r :

$$C_i(x) = \|\mathcal{F}_i(\pi(T_i \circ \pi^{-1}(x, D(x)))) - \mathcal{F}_r(x)\|_2, \quad (3)$$

where $\pi(\cdot)$ is the projection of 3D points in 3D space onto the image plane and $\pi^{-1}(x, D(x))$ is the inverse projection. The transformation converts 3D points from the camera space of I_r to that of I_i . The multiple cost values are averaged at each pixel x for multiple neighboring images.

Depth-pose Context-aware Temporal Attention (CTA)

The construction of the cost volume heavily relies on the static scene assumption, where it supposes that the object points remain static at time t and t^* . Thus, we re-project the features at time t to another plane with pose t^* at time t^* , to achieve the matching cost values. However, moving objects break this assumption since targets such as cars, trains, or pedestrians with a certain speed could move within the time gap. This gives rise to the feature inconsistency deviation, degraded (mismatching) cost values and re-projection loss, and finally drawbacks our optimization for depth and pose. We discard explicit settings such as the object motion prediction module or disentangle module [Lee *et al.*, 2021a; Feng *et al.*, 2022; Wimbauer *et al.*, 2021], which brings additional complexity and ignores the potential of complementary context-temporal information. Instead, we deliver our context-aware temporal attention (CTA) to implicitly rectify the mismatching problem, which efficiently cooperates the temporal features with context features via attention-based communication to achieve the feature consistency and estimation integrity.

Specifically, as shown in Fig. 4, taking depth optimization as an example, we first lift the temporal feature F_r^t to value (V) vectors via the mapping function $\sigma(\cdot)$. Meanwhile, we create query (Q) and key (K) vectors by adding the mapping functions $\theta(\cdot)$ and $\phi(\cdot)$ from the context feature F_r^c , and prepare long-range geometry embedding (LGE). We first allocate the query, key and value as $Q = \theta(F_r^c) \oplus LGE$, $K = \phi(F_r^c) \oplus LGE$, and $V = \sigma(F_r^t)$, respectively. Subsequently, the depth context-aware temporal attention is denoted as:

$$F_r^d = f_s(Q \otimes K) \otimes V \oplus F_r^t, \quad (4)$$

where f_s denotes the softmax operation, \oplus denotes the point-wise addition and \otimes denotes matrix multiplication.

Intuitively, compared with directly feeding F_r^t and F_r^c for refinement, our CTA explicitly aligns the features for moving objects through the cross-attention mechanism, to compensate for the mismatching discrepancy, which guarantees temporal-context feature fusion and seamless depth refinement. The temporal feature also helps the context feature to fulfill the moving target’s integrity, such as the second row of Fig. 1, where we rectify the wrong estimation in the ‘hole’ of the car head with temporal-content interaction. Similarly, for the pose optimization with fixed depth optimization, we employ context features F_i^c to extract query (Q) and key (K) vectors with added LGE. Hence, we allocate the query, key and value as $Q = \theta(F_i^c) \oplus LGE$, $K = \phi(F_i^c) \oplus LGE$, and $V = \sigma(F_i^t)$, respectively. Particularly, we lift the temporal feature F_i^t for value (V) vectors and finally adopt similar attention for pose context-aware temporal attention:

$$F_i^p = f_s(Q \otimes K) \otimes V \oplus F_i^t. \quad (5)$$

3.4 Long-range Geometry Embedding (LGE)

Existing works [Gu *et al.*, 2023; Feng *et al.*, 2022] usually adopt the two-frame correlation, such as the cost volume constructed from two frames, which limits the temporal interest field and wastes temporal helpful clues for geometry reasoning within the originally long-range multiple frames.

Therefore, to expand the temporal geometry relation range, especially for moving targets within long-range frames, we present our long-range geometry embedding module to create long-range geometry embedding (LGE), which seizes the beneficial long-range temporal hints for geometry and is inserted into the depth-pose CTA to promote the present depth and pose refinement.

Specifically, we demonstrate the detailed structure in Fig. 2, where we build the long-range cost map with the reference frame I_r and another arbitrary frame $I_j (j \neq i)$. Hence, we can achieve N-1 cost maps with N-1 different-range frames. Notably, we achieve the temporal feature with long-range cost maps and convolution, which reveals the same as the procedure in CTA-Refiner. Afterwards, we propose the 1×1 Conv, ReLU and Conv to deliver the geometry embedding $\varphi(F_j)$. Specifically, the number of $\varphi(F_j), j \neq i$ is N-1. Then, we aggregate this long-range temporal knowledge to create our embedding F_{ge} via the addition operation:

$$F_{ge} = \sum_{j, j \neq i}^N \varphi(F_j), \quad (6)$$

which will be added to the query and key vectors within the depth-pose CTA. Our approach efficiently reveals most of the geometry clues within the long-range temporal frames. It is also regarded as a temporal geometry prior to enhance the temporal-context association process.

3.5 Supervised Training Loss

We train our network by optimizing both depth and pose errors. Then, we formulate the depth loss as the $L1$ distance between the predicted depth map D and the ground truth \hat{D} :

$$\mathcal{L}_{\text{depth}} = \sum_{s=1}^m \gamma^{m-s} \|D^s - \hat{D}\|_1, \quad (7)$$

where the discounting factor γ is 0.85 and s denotes the stage number. There are m alternating update stages of depth and pose refinements. At each stage, we repeatedly refine the depth and pose n times. Next, our pose loss is defined based on the ground truth depth \hat{D} and pose \hat{T}_i with I_i relative to the reference image I_r :

$$\mathcal{L}_{\text{pose}} = \sum_{s=1}^m \sum_x \gamma^{m-s} \|\pi(T_i^s \circ \pi^{-1}(x, \hat{D}(x))) - \pi(\hat{T}_i \circ \pi^{-1}(x, \hat{D}(x)))\|_1, \quad (8)$$

where \circ means the Hadamard product. The pose loss summarizes the re-projection deviation of the pixel x according to the estimated camera pose T_i^s and the true pose \hat{T}_i in each stage. $\pi(\cdot)$ is the projection of 3D points in 3D space onto the image plane. Its inverse projection $\pi^{-1}(x, \hat{D}(x))$ maps the pixel x and its ground truth depth $\hat{D}(x)$ back points in the 3D space. Finally, the total supervised loss is calculated by:

$$\mathcal{L}_{\text{supervised}} = \mathcal{L}_{\text{depth}} + \mathcal{L}_{\text{pose}}. \quad (9)$$

4 Experiments and Results

4.1 Datasets

KITTI. KITTI dataset [Geiger *et al.*, 2012] is a popular benchmark for the task of autonomous driving, which provides over 93,000 depth maps with corresponding raw LiDAR scans and RGB images aligned with raw data. In experiments, we follow the widely-used KITTI Eigen split [Eigen *et al.*, 2014] for network training, which is composed of 22,600 images from 32 scenes for training and 697 images from 29 scenes for testing. The corresponding depth of each RGB image is sampled sparsely by the LiDAR sensor.

Virtual KITTI 2. VKITTI2 dataset [Gaidon *et al.*, 2016] is widely used for video understanding tasks, which consists of 5 sequence clones from the KITTI tracking benchmark and contains 50 high-resolution monocular videos generated from five different virtual worlds in urban settings under various imaging and weather conditions. These photo-realistic synthetic videos are fully annotated with depth labels.

NuScenes. NuScenes dataset [Caesar *et al.*, 2020] is a large-scale multi-modal autonomous driving dataset that is the first to carry the completely autonomous vehicle sensor suite: 32-beam LiDAR, 6 cameras and 5 radars with 360° coverage. It comprises 1,000 scenes, where each scene lasts 20 seconds and is fully annotated with 3D bounding boxes for 23 classes and 8 attributes.

4.2 Implementation Details

We implement our CTA-Depth in PyTorch and train it for 100 epochs with a mini-batch size of 4. The learning rate is 2×10^{-4} for both depth and pose refinement, which is decayed by a constant step (gamma=0.5 and step size=30). We set $\beta_1 = 0.9$ and $\beta_2 = 0.999$ in the Adam optimizer. We resize the input images to 320×960 for training, and set the number of sequential images to 2 for CTA-Refiner by balancing both computation efficiency and prediction accuracy. For long-range geometry embedding, the number of temporally adjacent images is set to $N = 3$. Since the output $\varphi(F_j)$ of the LGE for the same image is fixed and not updated with the iterations, this provides more prior temporal information to CTA-Refiner while ensuring network efficiency. We fix m at 3 and n at 4 in experiments.

4.3 Computation Time Analysis

Given the same Nvidia RTX A6000 GPU on the KITTI dataset, compared to the state-of-the-art one-frame method NeWCRFs [Yuan *et al.*, 2022], the inference speed of our CTA-Depth, i.e., the number of images inferred per second (FPS), is greatly improved by 58.9%, i.e., 5.53 (Ours) vs. 3.48 (NeWCRFs). This is because NeWCRFs use four swin-large transformers as multi-level encoders, while our method uses the lightweight ResNet18 backbone as the encoder to extract features. On the other hand, as shown in Table 2, although the FPS of CTA-Depth is slightly lower than that of the multi-frame method DRO [Gu *et al.*, 2023] (5.53 vs. 6.25), our performance significantly outperforms DRO and MaGNet [Bae *et al.*, 2022].

Method	Input	GT type	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE _{log} ↓	$\delta_1 < 1.25 \uparrow$	$\delta_2 < 1.25^2 \uparrow$	$\delta_3 < 1.25^3 \uparrow$
PackNet-SfM [Guizilini <i>et al.</i> , 2020]	M→S	Velodyne	0.090	0.618	4.220	0.179	0.893	0.962	0.983
DRO [Gu <i>et al.</i> , 2023]	M	Velodyne	0.073	0.528	3.888	0.163	0.924	0.969	0.984
CTA-Depth (ours)	M	Velodyne	0.071	0.496	3.598	0.143	0.931	0.975	0.989
BTS [Lee <i>et al.</i> , 2019]	S	Improved	0.059	0.241	2.756	0.096	0.956	0.993	0.998
GLPDepth [Kim <i>et al.</i> , 2022]	S	Improved	0.057	–	2.297	0.086	0.967	0.996	0.999
PackNet-SfM [Guizilini <i>et al.</i> , 2020]	M→S	Improved	0.064	0.300	3.089	0.108	0.943	0.989	0.997
BANet [Tang and Tan, 2018]	M	Improved	0.083	–	3.640	0.134	–	–	–
DeepV2D(2-view) [Teed and Deng, 2019]	M	Improved	0.064	0.350	2.946	0.120	0.946	0.982	0.991
DRO [Gu <i>et al.</i> , 2023]	M	Improved	0.047	0.199	2.629	0.082	0.970	0.994	0.998
CTA-Depth	M	Improved	0.038	0.145	2.224	0.069	0.978	0.996	0.999

Table 1: Quantitative results of supervised monocular depth estimation methods on the KITTI Eigen split. Note that the seven widely-used metrics are calculated strictly following the baseline [Gu *et al.*, 2023] and ground-truth median scaling is applied. “M” and “S” mean multiple and single frame. “M→S” means monocular multiple frame images are used in training while only a single frame image is used for inference. We utilize bold to highlight the best results and color the multi-frame baseline results and our performance gain over them in blue.

Method	Reference	Input	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE _{log} ↓	$\delta_1 < 1.25 \uparrow$	$\delta_2 < 1.25^2 \uparrow$	$\delta_3 < 1.25^3 \uparrow$	FPS ↑
Xu <i>et al.</i> [Xu <i>et al.</i> , 2018]	CVPR 2018	S	0.122	0.897	4.677	–	0.818	0.954	0.985	–
DORN [Fu <i>et al.</i> , 2018]	CVPR 2018	S	0.072	0.307	2.727	0.120	0.932	0.984	0.995	–
Yin <i>et al.</i> [Yin <i>et al.</i> , 2019]	ICCV 2019	S	0.072	–	3.258	0.117	0.938	0.990	0.998	–
PackNet-SAN [Guizilini <i>et al.</i> , 2021]	CVPR 2021	S	0.062	–	2.888	–	0.955	–	–	–
DPT* [Ranftl <i>et al.</i> , 2021]	ICCV 2021	S	0.062	–	2.573	0.092	0.959	0.995	0.999	–
PWA [Lee <i>et al.</i> , 2021b]	AAAI 2021	S	0.060	0.221	2.604	0.093	0.958	0.994	0.999	–
AdaBins [Bhat <i>et al.</i> , 2021]	CVPR 2021	S	0.058	0.190	2.360	0.088	0.964	0.995	0.999	2.96
NeWCRFs [Yuan <i>et al.</i> , 2022]	CVPR 2022	S	0.052	0.155	2.129	0.079	0.974	0.997	0.999	3.48
P3Depth [Patil <i>et al.</i> , 2022]	CVPR 2022	S	0.071	0.270	2.842	0.103	0.953	0.993	0.998	–
BANet [Tang and Tan, 2018]	ICLR 2019	M	0.083	–	3.640	0.134	–	–	–	–
MaGNet [Bae <i>et al.</i> , 2022]	CVPR 2022	M	0.054	0.162	2.158	0.083	0.971	–	–	–
DRO [Gu <i>et al.</i> , 2023]	RA-L 2023	M	0.059	0.230	2.799	0.092	0.964	0.994	0.998	6.25
CTA-Depth	–	M	0.045	0.156	2.275	0.073	0.978	0.997	0.999	5.53

Table 2: Quantitative results on KITTI Eigen split with the cap of 0-80m. Note that the seven widely-used metrics are calculated strictly following AdaBins [Bhat *et al.*, 2021]. “Abs Rel” error occupies the main ranking metric. “*” means using additional data for training. We utilize bold to highlight the best results of single-frame methods and multi-frame methods.

4.4 Evaluation of Our Method

Evaluation on KITTI. We first compare our CTA-Depth against top performers of supervised monocular depth estimation on the KITTI dataset; see Tables 1 & 2 for the results. For a fair comparison, all methods are evaluated given the same sequential images. In Table 1, the seven widely-used evaluation metrics are calculated strictly following the work [Gu *et al.*, 2023] and the ground-truth median scaling is applied to obtain the final output. In contrast, the seven evaluation metrics in Table 2 are calculated according to Adabins [Bhat *et al.*, 2021] and the final output is calculated by taking the average prediction of the image and its mirror image. Clearly, CTA-Depth achieves state-of-the-art performance over all the evaluation metrics under two different evaluation strategies. We further show the qualitative comparisons in Fig. 5 by comparing our method (c) with the recent approach [Gu *et al.*, 2023] (b). As shown in the green boxes, our method yields finer depth estimation results for moving objects in dynamic scenes, small objects and object contours, such as the trams and traffic signs. In addition, as shown in the top row of Fig. 5, our predicted depth map for the tram window is more consistent with the rest parts of the tram.

Evaluation on Virtual KITTI 2. We further verify our method on the virtual KITTI 2 dataset as shown in Table 3. We use a subset of the virtual KITTI 2, which contains 1,700 images for training and 193 images for testing. Notably,

Method	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE _{log} ↓
GLPDepth	0.058	0.217	2.146	0.125
Adabins	0.041	0.164	1.981	0.094
DRO	0.040	0.153	1.903	0.092
CTA-Depth	0.035	0.129	1.715	0.085

Table 3: Quantitative results on a subset of the Virtual KITTI 2.

Method	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE _{log} ↓
GLPDepth	0.061	0.340	3.159	0.121
Adabins	0.058	0.314	3.156	0.117
DRO	0.057	0.303	3.004	0.114
CTA-Depth	0.050	0.252	2.786	0.104

Table 4: Quantitative results on a subset of the nuScenes dataset.

our CTA-Depth achieved significantly better results than the multi-frame baseline methods over all evaluation metrics.

Evaluation on nuScenes. To further demonstrate the competitiveness of our approach, we also conduct an evaluation on the nuScenes dataset. In this experiment, we manually split a subset consisting of 2,600 images for training and 170 images for testing. The result is shown in Table 4. Again, the results show that our proposed method outperforms the baselines with a significant margin in all evaluation metrics.

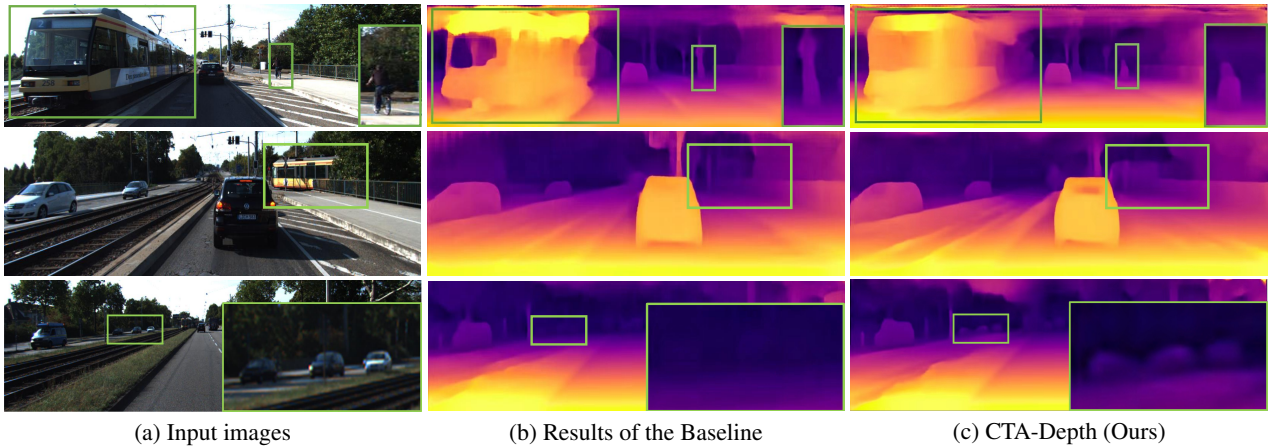


Figure 5: Qualitative comparisons given input images (a) from KITTI. Clearly, our method (c) yields more accurate depth maps than the Baseline method; (b) see particularly the regions marked by green boxes.

Setting	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE _{log} ↓
Baseline	0.060	0.275	3.132	0.104
+ Multi-level	0.059	0.267	3.105	0.101
+ Cross-scale	0.057	0.248	2.951	0.096
+ MAE	0.056	0.243	2.930	0.094
+ Depth-CTA	0.051	0.192	2.586	0.082
+ Pose-CTA	0.047	0.169	2.307	0.075
+ LGE	0.045	0.156	2.275	0.073
I_{t-1}, I_t	0.050	0.221	2.813	0.089
I_{t-2}, I_{t-1}, I_t	0.048	0.215	2.807	0.086
I_{t-1}, I_t, I_{t+1}	0.045	0.156	2.275	0.073

Table 5: Ablation study on the KITTI dataset. “Multi-level” refers to multi-level feature extraction, “Cross-scale” refers to cross-scale attention layers. “ I_{t-i} ” is the input frame at time $t - i$.

4.5 Ablation Study

To inspect the importance of each module in our method, we conduct an ablation study on the KITTI dataset and provide the results in Table 5. From top to bottom, the proposed modules are added in turn until the full method is constructed.

Baseline. To verify the effectiveness of each component, we build a baseline model. This model has a similar network architecture as the full pipeline, which includes the encoder-decoder structure with a deep recurrent network. In other words, the proposed CTA module and LGM are removed. The MAE module keeps only the single-level ResNet18 feature net for the depth and pose estimations.

Multi-level feature extraction. We extract four levels of feature maps as shown in Fig. 3. The performance gain shows that it provides more valuable information for the model.

Cross-scale attention. Next, we use the same ResNet18 as the feature encoder and add cross-scale attention layers to decode the features at each level following Fig. 3.

MAE module. In addition to the two operations above, the rearrange upscale is added to reduce the network complexity. Together with the PPM, they consequently enhance the estimation performance.

Depth-CTA and Pose-CTA. We add Depth-CTA and Pose-CTA after the cost map to obtain dynamic features by implicitly modeling the temporal relation and learn the geometric constraints through the temporal frames. In this way, the learned adaptive dynamic features are fed to the GRU optimizers and yield a noticeable performance gain. The “Abs Rel” error is reduced from 0.056 to 0.047.

LGE. The long-range geometry embedding module provides temporal priors for dynamic objects in several temporal neighboring frames and enhances the learning of the CTA module by large margins.

Multi-frame input. Here, we set up three experiments with different numbers of input frames. The results show that the optimal performance is achieved when the adopted temporal frames are I_{t-1}, I_t and I_{t+1} , i.e., two image pairs: (I_{t-1}, I_t) and (I_t, I_{t+1}) . However, the network becomes over-complicated and time-consuming to train or infer when utilizing more than three sample frames for depth refinement.

5 Conclusion

In this work, we present a novel CTA-Depth for multi-frame monocular depth estimation. To resolve the ambiguity caused by challenging dynamic scenes with moving objects, we propose the CTA-Refiner by designing context-aware temporal attention to implicitly leverage temporal information and model image self-similarities. In addition, we develop a novel long-range geometry embedding module to efficiently inject our refiner with geometry reasoning among the long-range temporal frames. Furthermore, we build a multi-level encoder-decoder network with the attention-enhanced predictor to obtain features with both global and local attentions. We achieve state-of-the-art performances on three challenging monocular depth estimation benchmarks. In the future, we would like to employ our multi-frame association mechanism in relevant tasks such as 3D object detection.

Contribution Statement

Zizhang Wu, Zhuozheng Li and Zhi-Gang Fan contributed equally to this work.

References

- [Bae *et al.*, 2022] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Multi-view depth estimation by fusing single-view depth probability with multi-view geometry. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2842–2851, 2022.
- [Bhat *et al.*, 2021] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. AdaBins: Depth estimation using adaptive bins. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4009–4018, 2021.
- [Caesar *et al.*, 2020] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 11621–11631, 2020.
- [Dosovitskiy *et al.*, 2021] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Int. Conf. on Learning Representations (ICLR)*. OpenReview.net, 2021.
- [Eigen *et al.*, 2014] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Conference and Workshop on Neural Information Processing Systems (NeurIPS)*, pages 2366–2347, 2014.
- [Feng *et al.*, 2022] Ziyue Feng, Liang Yang, Longlong Jing, Haiyan Wang, YingLi Tian, and Bing Li. Disentangling object motion and occlusion for unsupervised multi-frame monocular depth. In *European Conf. on Computer Vision (ECCV)*, pages 228–244, 2022.
- [Fu *et al.*, 2018] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2002–2011, 2018.
- [Gaidon *et al.*, 2016] Adrien Gaidon, Qiao Wang, Johann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4340–4349, 2016.
- [Geiger *et al.*, 2012] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361, 2012.
- [Gu *et al.*, 2023] Xiaodong Gu, Weihao Yuan, Zuozhuo Dai, Siyu Zhu, Chengzhou Tang, Zilong Dong, and Ping Tan. Dro: Deep recurrent optimizer for video to depth. *IEEE Robotics and Automation Letters*, 8(5):2844–2851, 2023.
- [Guizilini *et al.*, 2020] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2485–2494, 2020.
- [Guizilini *et al.*, 2021] Vitor Guizilini, Rares Ambrus, Wolfram Burgard, and Adrien Gaidon. Sparse auxiliary networks for unified monocular depth prediction and completion. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 11078–11088, 2021.
- [Kim *et al.*, 2022] Doyeon Kim, Woonghyun Ga, Pyunghwan Ahn, Donggyu Joo, Sehwan Chun, and Junmo Kim. Global-local path networks for monocular depth estimation with vertical cutdepth. *arXiv preprint arXiv:2201.07436*, 2022.
- [Klingner *et al.*, 2020] Marvin Klingner, Jan-Aike Termöhlen, Jonas Mikolajczyk, and Tim Fingscheidt. Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In *European Conf. on Computer Vision (ECCV)*, pages 582–600. Springer, 2020.
- [Laina *et al.*, 2016] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *Int. Conf. on 3D Vision (3DV)*, pages 239–248, 2016.
- [Lee and Kim, 2020] Jae-Han Lee and Chang-Su Kim. Multi-loss rebalancing algorithm for monocular depth estimation. In *European Conf. on Computer Vision (ECCV)*, pages 785–801, 2020.
- [Lee *et al.*, 2019] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019.
- [Lee *et al.*, 2021a] Seokju Lee, Sunghoon Im, Stephen Lin, and In So Kweon. Learning monocular depth in dynamic scenes via instance-aware projection consistency. In *AAAI Conf. on Artificial Intell. (AAAI)*, pages 1863–1872, 2021.
- [Lee *et al.*, 2021b] Sihaeng Lee, Janghyeon Lee, Byungju Kim, Eojindl Yi, and Junmo Kim. Patch-wise attention network for monocular depth estimation. In *AAAI Conf. on Artificial Intell. (AAAI)*, pages 1873–1881, 2021.
- [Li *et al.*, 2015] Bo Li, Chunhua Shen, Yuchao Dai, Anton Van Den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1119–1127, 2015.
- [Li *et al.*, 2019] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T Freeman. Learning the depths of moving people by watching frozen people. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4521–4530, 2019.
- [Li *et al.*, 2022] Yingyan Li, Yuntao Chen, Jiawei He, and Zhaoxiang Zhang. Densely constrained depth estimator for monocular 3d object detection. In *European Conf. on Computer Vision (ECCV)*, pages 718–734. Springer, 2022.

- [Lin *et al.*, 2017] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2117–2125, 2017.
- [Liu *et al.*, 2015] Fayao Liu, Chunhua Shen, and Guosheng Lin. Deep convolutional neural fields for depth estimation from a single image. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5162–5170, 2015.
- [Liu *et al.*, 2018] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8759–8768, 2018.
- [Liu *et al.*, 2021] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *IEEE Int. Conf. on Computer Vision (ICCV)*, pages 10012–10022, 2021.
- [Mumuni and Mumuni, 2022] Fuseini Mumuni and Alhasan Mumuni. Bayesian cue integration of structure from motion and cnn-based monocular depth estimation for autonomous robot navigation. *International Journal of Intelligent Robotics and Applications*, 6(2):191–206, 2022.
- [Patil *et al.*, 2020] 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.
- [Patil *et al.*, 2022] Vaishakh Patil, Christos Sakaridis, Alexander Liniger, and Luc Van Gool. P3Depth: Monocular depth estimation with a piecewise planarity prior. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1610–1621, 2022.
- [Ranftl *et al.*, 2021] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *IEEE Int. Conf. on Computer Vision (ICCV)*, pages 12179–12188, 2021.
- [Ricci *et al.*, 2018] Elisa Ricci, Wanli Ouyang, Xiaogang Wang, Nicu Sebe, et al. Monocular depth estimation using multi-scale continuous crfs as sequential deep networks. *IEEE Trans. Pattern Anal. & Mach. Intell.*, 41(6):1426–1440, 2018.
- [Tang and Tan, 2018] Chengzhou Tang and Ping Tan. BANet: Dense bundle adjustment networks. In *Int. Conf. on Learning Representations (ICLR)*, 2018.
- [Teed and Deng, 2019] Zachary Teed and Jia Deng. DeepV2D: Video to depth with differentiable structure from motion. In *Int. Conf. on Learning Representations (ICLR)*, 2019.
- [Teed and Deng, 2020] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *European Conf. on Computer Vision (ECCV)*, pages 402–419, 2020.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Lion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Conference and Workshop on Neural Information Processing Systems (NeurIPS)*, pages 6000–6010, 2017.
- [Wang *et al.*, 2015] Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, and Alan L Yuille. Towards unified depth and semantic prediction from a single image. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2800–2809, 2015.
- [Wang *et al.*, 2019] Rui Wang, Stephen M Pizer, and Jan-Michael Frahm. Recurrent neural network for (un-) supervised learning of monocular video visual odometry and depth. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5555–5564, 2019.
- [Watson *et al.*, 2021] Jamie Watson, Oisín Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1164–1174, 2021.
- [Wimbauer *et al.*, 2021] Felix Wimbauer, Nan Yang, Lukas Von Stumberg, Niclas Zeller, and Daniel Cremers. MonoRec: Semi-supervised dense reconstruction in dynamic environments from a single moving camera. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 6112–6122, 2021.
- [Wu *et al.*, 2022] Zizhang Wu, Zhi-Gang Fan, Zhuozheng Li, Jizheng Wang, Tianhao Xu, Qiang Tang, Fan Wang, and Zhengbo Luo. Monocular fisheye depth estimation for automated valet parking: Dataset, baseline and deep optimizers. In *International Conference on Intelligent Transportation Systems (ITSC)*, pages 01–07. IEEE, 2022.
- [Xu *et al.*, 2018] Dan Xu, Wei Wang, Hao Tang, Hong Liu, Nicu Sebe, and Elisa Ricci. Structured attention guided convolutional neural fields for monocular depth estimation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3917–3925, 2018.
- [Yang *et al.*, 2021] Guanglei Yang, Hao Tang, Mingli Ding, Nicu Sebe, and Elisa Ricci. Transformers solve the limited receptive field for monocular depth prediction. *arXiv preprint arXiv:2103.12091*, 2021.
- [Yin *et al.*, 2019] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *IEEE Int. Conf. on Computer Vision (ICCV)*, pages 5684–5693, 2019.
- [Yuan *et al.*, 2022] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. NeWCRFs: Neural window fully-connected crfs for monocular depth estimation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3916–3925, 2022.
- [Zhang *et al.*, 2019] Haokui Zhang, Chunhua Shen, Ying Li, Yuanzhouhan Cao, Yu Liu, and Youliang Yan. Exploiting temporal consistency for real-time video depth estimation. In *IEEE Int. Conf. on Computer Vision (ICCV)*, pages 1725–1734, 2019.
- [Zhao *et al.*, 2017] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2881–2890, 2017.