

This CVPR 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

MonoRec: Semi-Supervised Dense Reconstruction in Dynamic Environments from a Single Moving Camera

Felix Wimbauer^{1,*} Nan Yang^{1,2,*} Lukas von Stumberg¹ Niclas Zeller^{1,2} Daniel Cremers^{1,2} ¹ Technical University of Munich, ² Artisense

{wimbauer, yangn, stumberg, zellern, cremers}@in.tum.de

Abstract

In this paper, we propose MonoRec, a semi-supervised monocular dense reconstruction architecture that predicts depth maps from a single moving camera in dynamic environments. MonoRec is based on a multi-view stereo setting which encodes the information of multiple consecutive images in a cost volume. To deal with dynamic objects in the scene, we introduce a MaskModule that predicts moving object masks by leveraging the photometric inconsistencies encoded in the cost volumes. Unlike other multi-view stereo methods, MonoRec is able to reconstruct both static and moving objects by leveraging the predicted masks. Furthermore, we present a novel multi-stage training scheme with a semi-supervised loss formulation that does not require LiDAR depth values. We carefully evaluate MonoRec on the KITTI dataset and show that it achieves state-of-theart performance compared to both multi-view and singleview methods. With the model trained on KITTI, we furthermore demonstrate that MonoRec is able to generalize well to both the Oxford RobotCar dataset and the more challenging TUM-Mono dataset recorded by a handheld camera. Code and related materials are available at https: //vision.in.tum.de/research/monorec.

1. Introduction

1.1. Real-world Scene Capture from Video

Obtaining a 3D understanding of the entire static and dynamic environment can be seen as one of the key-challenges in robotics, AR/VR, and autonomous driving. State of today, this is achieved based on the fusion of multiple sensor sources (incl. cameras, LiDARs, RADARs and IMUs). This guarantees dense coverage of the vehicle's surroundings and accurate ego-motion estimation. However, driven by the high cost as well as the challenge to maintain crosscalibration of such a complex sensor suite, there is an in-



Figure 1: MonoRec can deliver high-quality dense reconstruction from a single moving camera. The figure shows an example of a large-scale outdoor point cloud reconstruction (KITTI Odometry sequence 07) by simply accumulating predicted depth maps. Please refer to our project page for the video of the entire reconstruction of the sequence.

creasing demand of reducing the total number of sensors. Over the past years, researchers have therefore put a lot of effort into solving the problem of perception with only a single monocular camera. Considering recent achievements in monocular visual odometry (VO) [8, 58, 51], with respect to ego-motion estimation, this was certainly successful. Nevertheless, reliable dense 3D mapping of the static environment and moving objects is still an open research topic.

To tackle the problem of dense 3D reconstruction based on a single moving camera, there are basically two paral-

^{*} Indicates equal contribution.

lel lines of research. On one side, there are dense multiview stereo (MVS) methods, which evolved over the last decade [39, 45, 2] and saw a great improvement through the use of convolutional neural networks (CNNs) [23, 61, 57]. On the other side, there are monocular depth prediction methods which purely rely on deep learning [7, 16, 58]. Though all these methods show impressive performance. both types have also their respective shortcomings. For MVS the overall assumption is a stationary environment to be reconstructed, so the presence of dynamic objects deteriorate their performance. Monocular depth prediction methods, in contrast, perform very well in reconstructing moving objects, as predictions are made only based on individual images. At the same time, due to their use of a single image only, they strongly rely on the perspective appearance of objects as observed with specific camera intrinsics and extrinsics and therefore do not generalize well to other datasets.

1.2. Contribution

To combine the advantage of both deep MVS and monocular depth prediction, we propose MonoRec, a novel monocular dense reconstruction architecture that consists of a MaskModule and a DepthModule. We encode the information from multiple consecutive images using cost volumes which are constructed based on structural similarity index measure (SSIM) [54] instead of sum of absolute differences (SAD) like prior works. The MaskModule is able to identify moving pixels and downweights the corresponding voxels in the cost volume. Thereby, in contrast to other MVS methods, MonoRec does not suffer from artifacts on moving objects and therefore delivers depth estimations on both static and dynamic objects.

With the proposed multi-stage training scheme, MonoRec achieves state-of-the-art performance compared to other MVS and monocular depth prediction methods on the KITTI dataset [14]. Furthermore, we validate the generalization capabilities of our network on the Oxford RobotCar dataset [35] and the TUM-Mono dataset [9]. Figure 1 shows a dense point cloud reconstructed by our method on one of our test sequences of KITTI.

2. Related Work

2.1. Multi-view Stereo

Multi-view stereo (MVS) methods estimate a dense representation of the 3D environment based on a set of images with known poses. Over the past years, several methods have been developed to solve the MVS problem [46, 28, 30, 2, 47, 49, 39, 13, 45, 60] based on classical optimization. Recently, due to the advance of deep neural networks (DNNs), different learning based approaches were proposed. This representation can be volumetric

[26, 27, 36] or 3D point cloud based [3, 12]. Most popular are still depth map representations predicted from a 3D cost volume [23, 53, 61, 66, 22, 56, 41, 24, 33, 62, 19, 64, 57]. Huang et al. [23] proposed one of the first cost-volume based approaches. They compute a set of image-pair-wise plane-sweep volumes with respect to a reference image and use a CNN to predict one single depth map based on this set. Zhou et al. [66] also use the photometric cost volumes as the inputs of the deep neural networks and employ a two stage approach for dense depth prediction. Yao et al. [61] instead calculate a single cost volume using deep features of all input images.

2.2. Dense Depth Estimation in Dynamic Scenes

Reconstructing dynamic scenes is challenging since the moving objects violate the static-world assumption for classical multi-view stereo methods. Russell et al. [43] and Ranftl et al. [40] base on motion segmentation and perform classical optimization. Li et al. [32] proposed to estimate dense depth maps from the scenes with moving people. All these methods need additional inputs, e.g., optical flow, object masks, etc., for the inference, while MonoRec requires only the posed images as the inputs. Another line of research is monocular depth estimation [7, 6, 29, 31, 11, 59, 16, 48, 67, 63, 65, 52, 18, 17, 58]. These methods are not affected by moving objects, but the depth estimation is not necessarily accurate, especially in unseen scenarios. Luo et al. [34] proposed a test-time optimization method which is not real-time capable. In a concurrent work, Watson et al. [55] address moving objects with the consistency between monocular depth estimation and multi-view stereo, while MonoRec predicts the dynamic masks explicitly by the proposed MaskModule.

2.3. Dense SLAM

Several of the methods cited above solve both the problem of dense 3D reconstruction and camera pose estimation [48, 67, 63, 65, 66, 59, 58]. Nevertheless, these methods either solve both problems independently or only integrate one into the other (e.g. [66, 58]). Newcombe et al. [37] instead jointly optimize the 6DoF camera pose and the dense 3D scene structure. However, due to its volumetric map representation it is only applicable to smallscale scenes. Recently, Bloesch et al. [1] proposed a learned code representation which can be optimized jointly with the 6DoF camera poses. This idea is pursued by Czarnowski et al. [5] and integrated into a full SLAM system. All the above-mentioned methods, however, do not address the issue of moving objects. Instead, the proposed MonoRec network explicitly deals with moving objects and achieves superior accuracy both on moving and on static structures. Furthermore, prior works show that the accuracy of camera tracking does not necessarily improve with more



Figure 2: **MonoRec Architecture**: It first constructs a photometric cost volume from multiple input frames. Unlike prior works, we use the SSIM [54] metric instead of SAD to measure the photometric consistency. The MaskModule aims to detect inconsistencies between the different input frames to determine moving objects. The multi-frame cost volume C is multiplied with the predicted mask and then passed to the DepthModule which predicts a dense inverse depth map. In both the decoders of MaskModule and DepthModule, the cost volume features are concatenated with pre-trained ResNet-18 features.

points [8, 10]. MonoRec therefore focuses solely on delivering dense reconstruction using poses from a sparse VO system and shows state-of-the-art results on public benchmarks. Note that, this way, MonoRec can be easily combined with any VO systems with arbitrary sensor setups.

3. The MonoRec Network

MonoRec uses a set of consecutive frames and the corresponding camera poses to predict a dense depth map for the given keyframe. The MonoRec architecture combines a MaskModule and a DepthModule. MaskModule predicts moving object masks that improve depth accuracy and allows us to eliminate noise in 3D reconstructions. Depth-Module predicts a depth map from the masked cost volume. In this section, we first describe the different modules of our architecture, and then discuss the specialized multi-stage semi-supervised training scheme.

3.1. Preliminaries

Our method aims to predict a dense inverse depth map D_t of the selected keyframe from a set of consecutive frames $\{I_1, \dots, I_N\}$. We denote the selected keyframe as I_t and others as $I_{t'}$ ($t' \in \{1, \dots, N\} \setminus t$). Given the camera intrinsics, the inverse depth map D_t , and the relative camera pose $\mathbf{T}_{t'}^t \in SE(3)$ between $I_{t'}$ and I_t , we can perform the reprojection from $I_{t'}$ to I_t as

$$I_{t'}^{t} = I_{t'} \left\langle proj\left(D_{t}, \mathbf{T}_{t'}^{t}\right) \right\rangle, \tag{1}$$

where proj() is the projection function and $\langle \rangle$ is the differentiable sampler [25]. This reprojection formulation is important for both the cost volume formation (Sec. 3.2) and the self-supervised loss term (Sec. 3.4).

In the following, we refer to the consecutive frames as temporal stereo (**T**) frames. During training, we use an additional static stereo (**S**) frame I_{t^S} for each sample, which was captured by a synchronized stereo camera at the same time as the respective keyframe.

3.2. Cost Volume

A cost volume encodes geometric information from the different frames in a tensor that is suited as input for neural networks. For a number of discrete depth steps, the temporal stereo frames are reprojected to the keyframe and a pixel-wise photometric error is computed. Ideally, the lower the photometric error, the better the depth step approximates the real depth at a given pixel. Our cost volume follows the general formulation of the prior works [37, 66]. Nevertheless, unlike the previous works that define the photometric error pe() as a patch-wise SAD, we propose to use the SSIM as follows:

$$pe(\mathbf{x}, d) = \frac{1 - \text{SSIM}(I_{t'}^t(\mathbf{x}, d), I_t(\mathbf{x}))}{2}$$
(2)

with 3×3 patch size. Here $I_{t'}^t(\mathbf{x}, d)$ defines the intensity at pixel \mathbf{x} of the image $I_{t'}$ warped with constant depth d. In practice, we clamp the error to [0, 1]. The cost volume C stores at $C(\mathbf{x}, d)$ the aggregated photometric consistency for pixel \mathbf{x} and depth d

$$C(\mathbf{x},d) = 1 - 2 \cdot \frac{1}{\sum_{t'} \omega_{t'}} \cdot \sum_{t'} p e_{t'}^t(\mathbf{x},d) \cdot \omega_{t'}(\mathbf{x}) \quad (3)$$

where $d \in \{d_i | d_{min} + \frac{i}{M} \cdot (d_{min} - d_{max})\}$. The weighting term $w_{t'}(\mathbf{x})$ weights the optimal depth step height based on the photometric error while others are weighted lower:

$$w_{t'}(\mathbf{x}) = 1 - \frac{1}{M-1}$$

$$\cdot \sum_{d \neq d^*} \exp\left(-\alpha \left(p e_{t'}^t(\mathbf{x}, d) - p e_{t'}^t(\mathbf{x}, d^*)\right)^2\right) \quad (4)$$

with $d_{t'}^* = \arg \min_d p e_{t'}^t(\mathbf{x}, d)$. Note that $C(\mathbf{x}, d)$ has the range [-1, 1] where -1/1 indicates the lowest/highest photometric consistency.

In the following section, we denote cost volumes calculated based on the keyframe I_t and only *one* non-keyframe $I_{t'}$ by $C_{t'}(\mathbf{x}, d)$ where applicable.

3.3. Network Architecture

As shown in Figure 2, the proposed network architecture contains two sub-modules, namely, MaskModule and DepthModule.

MaskModule MaskModule aims to predict a mask M_t where $M_t(\mathbf{x}) \in [0,1]$ indicates the probability of a pixel x in I_t belonging to a moving object. Determining moving objects from I_t alone is an ambiguous task and hard to be generalizable. Therefore, we propose to use the set of cost volumes $\{C_{t'}|t' \in \{1, \cdots, N\} \setminus t\}$ which encode the geometric priors between I_t and $\{I_{t'} | t' \in \{1, \cdots, N\} \setminus t\}$ respectively. We use $C_{t'}$ instead of C since the inconsistent geometric information from different $C_{t'}$ is a strong prior for moving object prediction - dynamic pixels yield inconsistent optimal depth steps in different $C_{t'}$. However, geometric priors alone are not enough to predict moving objects, since poorly-textured or non-Lambertian surfaces can lead to inconsistencies as well. Furthermore, the cost volumes tend to reach a consensus on wrong depths that semantically don't fit into the context of the scene for objects that move at constant speed. Therefore, we further leverage pre-trained ResNet-18 [21] features of I_t to encode semantic priors in addition to the geometric ones. The network adapts a U-Net architecture design [42] with skip connections. All cost volumes are passed through the encoders with shared weights. The features from different cost volumes are aggregated using max-pooling and then passed through the decoder. In this way, MaskModule can be applied to different numbers of frames without retraining.

DepthModule DepthModule predicts a dense pixel-wise inverse depth map D_t of I_t . To this end, the module receives the complete cost volume C concatenated with the keyframe I_t . Unlike MaskModule, here we use C instead of $C_{t'}$ since multi-frame cost volumes in general lead to higher depth accuracy and robustness against photometric noise [37]. To eliminate wrong depth predictions for moving objects, we perform pixel-wise multiplication between M_t and the cost volume C for every depth step d. This way, there won't be any maxima (*i.e.* strong priors) in regions of moving objects left, such that DepthModule has to rely on information from the image features and the surroundings to infer the depth of moving objects. We employ a U-Net architecture with multi-scale depth outputs from the decoder [17]. Finally, DepthModule outputs an interpolation factor between d_{min} and d_{max} . In practice, we use s = 4 scales of depth prediction.

3.4. Multi-stage Training

In this section, we propose a multi-stage training scheme for the networks. Specifically, the bootstrapping stage, the



Figure 3: **Auxiliary Training Masks**: Examples of auxiliary training masks from the training set that are used as reference.

MaskModule refinement stage and the DepthModule refinement stage are executed successively.

Bootstrapping In the bootstrapping stage, MaskModule and DepthModule are trained separately. DepthModule takes the *non-masked* C as the input and predicts D_t . The training objective of DepthModule is defined as a multiscale ($s \in [0, 3]$) semi-supervised loss. It combines a selfsupervised photometric loss and an edge-aware smoothness term, as proposed in [17], with a supervised sparse depth loss.

$$\mathcal{L}_{depth} = \sum_{s=0}^{3} \mathcal{L}_{self,s} + \alpha \mathcal{L}_{sparse,s} + \beta \mathcal{L}_{smooth,s}.$$
 (5)

The self-supervised loss is computed from the photometric errors between the keyframe and the reprojected temporal stereo and static stereo frames:

$$\mathcal{L}_{self,s} = \min_{t^{\star} \in t' \cup \{t^{S}\}} \left(\lambda \frac{1 - \text{SSIM}(I_{t^{\star}}^{t}, I_{t})}{2} + (1 - \lambda) ||I_{t^{\star}}^{t} - I_{t}||_{1} \right),$$

$$(6)$$

where $\lambda = 0.85$. Note that $\mathcal{L}_{self,s}$ takes the per-pixel minimum which has be shown to be superior compared to the per-pixel average [17]. The sparse supervised depth loss is defined as

$$\mathcal{L}_{sparse,s} = ||D_t - D_{VO}||_1, \tag{7}$$

where the ground-truth sparse depth maps (D_{VO}) are obtained by a visual odometry system [59]. Note that all the supervision signals of DepthModule are generated from either images themselves or the visual odometry system without any manual labeling or LiDAR depth.

MaskModule is trained with the mask loss \mathcal{L}_{mask} which is the weighted binary cross entropy between the predicted mask M_t and the auxiliary ground-truth moving object mask M_{aux} . We generate M_{aux} by leveraging a pre-trained Mask-RCNN and the trained DepthModule as explained above. We firstly define the movable object classes, e.g., cars, cyclists, etc, and then obtain the instance segmentations of these object classes for the training images. A movable instance is classified as a moving instance if it has a high ratio of photometrically inconsistent pixels between temporal stereo and static stereo. Specifically, for each image, we predict its depth maps D_t and D_t^S using the cost volumes formed by temporal stereo images C and static stereo images C^S , respectively. Then a pixel **x** is regarded as a moving pixel if two of the following three metrics are above predefined thresholds: (1) The static stereo photometric error using D_t , *i.e.*, $pe_{tS}^t(\mathbf{x}, D_t(\mathbf{x}))$. (2) The average temporal stereo photometric error using D_t^S , *i.e.*, $\overline{pe_{tS}^t}(\mathbf{x}, D_t^S(\mathbf{x}))$. (3) The difference between $D_t(\mathbf{x})$ and $D_t^S(\mathbf{x})$. Please refer to our supplementary materials for more details. Figure 3 shows some examples of the generated auxiliary ground-truth moving object masks.

MaskModule Refinement The bootstrapping stage for MaskModule is limited in two ways: (1) Heavy augmentation is needed since mostly only a very small percentage of pixels on the image belongs to moving objects. (2) The auxiliary masks are not necessarily related to the geometric prior in the cost volume, which slows down the convergence. Therefore, to improve the mask prediction, we utilize the trained DepthModule from the bootstrapping stage. We leverage the fact that the depth prediction for moving objects, and consequently the photometric consistency, should be better with a static stereo prediction than with a temporal stereo one. Therefore, similar to the classification of moving pixels as explained in the previous section, we obtain D_t^S and D_t from two forward passes using C^S and C as inputs, respectively. Then we compute the static stereo photometric error $L^{\prime S}_{self,s}$ using D^S_t as depth and the temporal stereo photometric error $L_{self,s}^{\prime T}$ using D_t as depth. To train M_t , we interpret it as pixel-wise interpolation factors between $\mathcal{L}_{self,s}^{\prime S}$ and $\mathcal{L}_{self,s}^{\prime T}$, and minimize the summation:

$$\mathcal{L}_{m_ref} = \sum_{s=0}^{3} \left(M_t \mathcal{L}_{depth,s}^{\prime S} + (1 - M_t) \mathcal{L}_{depth,s}^{\prime T} \right) + \mathcal{L}_{mask}.$$
(8)

Figure 4(a) shows the diagram illustrating different loss terms. Note that we still add the supervised mask loss \mathcal{L}_{mask} as a regularizer to stabilize the training. This way, the new gradients are directly related to the geometric structure in the cost volume and help to improve the mask prediction accuracy and alleviate the danger of overfitting.

DepthModule Refinement The bootstrapping stage does not distinguish between the moving pixels and static pixels when training DepthModule. Therefore, we aim to refine DepthModule such that it is able to predict proper depths also for moving objects. The key idea is that, by utilizing M_t , only the static stereo loss is backpropagated for moving pixels, while for static pixels the temporal stereo, static



Figure 4: **Refinement Losses:** a) MaskModule refinement and b) DepthModule refinement loss functions. Dashed outlines denote that no gradient is being computed for the respective forward pass in the module.

stereo and sparse depth losses are backpropagated. Because moving objects make up only a small percentage of all pixels in a keyframe, the gradients from the photometric error are rather weak. To solve this, we perform a further static stereo forward pass and use the resulting depth map D_t^S as prior for moving objects. Therefore, as shown in Figure 4(b), the loss for refining DepthModule is defined as

$$\mathcal{L}_{d.ref,s} = (1 - M_t) \left(\mathcal{L}_{self,s} + \alpha \mathcal{L}_{sparse,s} \right) + M_t \left(\mathcal{L}_{self,s}^S + \gamma \left| \left| D_t - D_t^S \right| \right|_1 \right)$$
(9)
+ $\beta \mathcal{L}_{smooth,s}.$

3.4.1 Implementation Details

The networks are implemented in PyTorch [38] with image size 512×256 . For the bootstrapping stage, we train Depth-Module for 70 epochs with learning rate $lr = 1e^{-4}$ for the first 65 epochs and $lr = 1e^{-5}$ for the remaining ones. MaskModule is trained for 60 epochs with $lr = 1e^{-4}$. During MaskModule refinement, we train for 32 epochs with $lr = 1e^{-4}$, and during DepthModule refinement we train for 15 epochs with $lr = 1e^{-4}$ and another 4 epochs at $lr = 1e^{-5}$. The hyperparameters α , β and γ are set to $4, 10^{-3} \times 2^{-s}$ and 4, respectively. For inference, MonoRec can achieve 10 fps with batch size 1 using 2GB memory.

4. Experiments

To evaluate the proposed method, we first compare against state-of-the-art monocular depth prediction and MVS methods with our train/test split of the KITTI dataset [15]. Then, we perform extensive ablation studies to show the efficacy of our design choices. In the end, we demonstrate the generalization capabilities of different methods on Oxford RobotCar [35] and TUM-Mono [9] using the model trained on KITTI.



Figure 5: **Qualitative Results on KITTI**: The upper part of the figure shows the results for a selected number of frames from the KITTI test set. The compared PackNet model was trained in a semi-supervised fashion using LiDAR as the ground truth. Besides the depth maps, we also show the 3D point clouds by reprojecting the depth and viewing from two different perspectives. For comparison we show the LiDAR ground truth from the corresponding perspectives. Our method clearly shows the best prediction quality. The lower part of the figure shows large scale reconstructions as point clouds accumulated from multiple frames. The red insets depict the reconstructed artifacts from moving objects. With the proposed MaskModule, we can effectively filter out the moving objects to avoid those artifacts in the final reconstruction.

Method	Training	Dataset	Input	Abs Rel	Sq Rel	RMSE	RMSE _{log}	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Colmap [44] (geometric)	-	-	KF + 2	0.099	3.451	5.632	0.184	0.952	0.979	0.986
Colmap [44] (photometric)	-	-	KF + 2	0.190	6.826	7.781	0.531	0.893	0.932	0.947
Monodepth2 [17]	MS	Eigen Split	KF	0.082	0.405	3.129	0.127	0.931	0.985	0.996
PackNet [20]	MS	CS+Eigen Split	KF	0.080	0.331	2.914	0.124	0.929	0.987	0.997
PackNet [20]	MS, D	CS+Eigen Split	KF	0.077	0.290	2.688	0.118	0.935	0.988	0.997
DORN [11]	D	Eigen Split	KF	0.077	0.290	2.723	0.113	0.949	0.988	<u>0.996</u>
DeepMVS [23]	D	Odom. Split	KF+2	0.103	1.160	3.968	0.166	0.896	0.947	0.978
DeepMVS [23] (pretr.)	D	Odom. Split	KF+2	0.088	0.644	3.191	0.146	0.914	0.955	0.982
DeepTAM [66] (only FB)	MS, D*	Odom. Split	KF+2	0.059	0.474	2.769	0.096	0.964	0.987	0.994
DeepTAM [66] (1x Ref.)	MS, D*	Odom. Split	KF+2	0.053	0.351	2.480	0.089	<u>0.971</u>	<u>0.990</u>	0.995
MonoRec	MS, D*	Odom. Split	KF+2	0.050	<u>0.295</u>	2.266	0.082	0.973	0.991	<u>0.996</u>

Table 1: Quantitative Results on KITTI: Comparison between MonoRec and other methods on our KITTI test set. The Dataset column shows the training dataset used by the corresponding method and please note that Eigen split is a *superset* of our odometry split. Best / Second best results are marked **bold** / <u>underlined</u>. The evaluation result shows that our method achieves overall the best performance. Legend: M: *Monocular images*, S: *Stereo images*, D: *GT depth*, D*: *Depths from DVSO*, KF: *Keyframe*, KF + 2: *Keyframe* + 2 mono frames, CS: *Cityscapes* [4], pretr.: *Pretrained network*, FB: *Fixed band module of DeepTAM*, Ref.: *Narrow band refinement module of DeepTAM*





4.1. The KITTI Dataset

The Eigen split [6] is the most popular training/test split for evaluating depth estimation on KITTI. We cannot make use of it directly since MonoRec requires temporally continuous images with estimated poses. Hence, we select our training/testing splits as the intersection between the KITTI Odometry benchmark and the Eigen split, which results in 13714/8634 samples for training/testing. We obtain the relative poses between the images from the monocular VO system DVSO [59]. During training, we also leverage the point clouds generated by DVSO as the sparse depth supervision signals. For training MaskModule we only use images that contain moving objects in the generated auxiliary masks, 2412 in total. For all the following evaluation results we use the improved ground truth [50] and cap depths at 80 m.

We first compare our method against the recent state of the art including an optimization based method (Colmap), self-supervised monocular methods (MonoDepth2 and PackNet), a semi-supervised monocular method using sparse LiDAR data (PackNet), a supervised monocular method (DORN) and MVS methods (DeepMVS and DeepTAM), shown in Table 1. Note that the training code of DeepTAM was not published, we therefore implemented it ourselves for training and testing using our split to deliver a fair comparison. Our method outperforms all the other methods with a notable margin despite relying on images only without using LiDAR ground truth for training.

This is also clearly reflected in the qualitative results shown in Figure 5. Compared with monocular depth estimation methods, our method delivers very sharp edges in the depth maps and can recover finer details. In comparison to the other MVS methods, it can better deal with moving objects, which is further illustrated in Figure 7.

A single depth map usually cannot really reflect the quality for large scale reconstruction. We therefore also visualize the accumulated points using the depth maps from multiple frames in lower part of Figure 5. We can see that our method can deliver very high quality reconstruction and, due to our MaskModule, is able to remove artifacts caused by moving objects. We urge readers to watch the supplementary video for more convincing comparisons.

Ablation Studies. We also investigated the contribution of the different components towards the method's performance. Table 2 shows quantitative results of our ablation studies, which confirm that all our proposed contributions improve the depth prediction over the baseline method. Furthermore, Figure 6 demonstrates the qualitative improvement achieved by MaskModule and refinement training.

4.2. Oxford RobotCar and TUM-Mono

To demonstrate the generalization capabilities of MonoRec, we test our KITTI model on the Oxford Robot-Car dataset and the TUM-Mono dataset. Oxford RobotCar is a street view dataset and shows a similar motion pattern and view perspective to KITTI. TUM-Mono, however, is recorded by a handheld monochrome camera, so it demon-

Model	SSIM	MaskModule	D. Ref.	M. Ref.	Abs Rel	Sq Rel	RMSE	RMSE _{log}	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Baseline					0.056	0.342	2.624	0.092	0.965	0.990	0.994
Baseline	\checkmark				0.054	0.346	2.444	0.088	0.970	0.989	<u>0.995</u>
MonoRec	\checkmark	\checkmark		\checkmark	0.054	0.306	2.372	0.087	0.970	0.990	0.995
MonoRec	\checkmark		\checkmark		0.051	0.346	2.361	0.085	0.972	0.990	0.995
MonoRec	\checkmark	\checkmark	\checkmark		0.052	0.302	2.303	0.087	0.969	0.990	0.995
MonoRec	\checkmark	\checkmark	\checkmark	\checkmark	0.050	0.295	2.266	0.082	0.973	0.991	0.996

Table 2: **Ablation Study**: Baseline consists of only DepthModule using the unmasked cost volume (CV). Baseline without SSIM uses a 5x5 patch that has same receptive field as SSIM. Using SSIM to form CV gives a significant improvement. For MonoRec, only the addition of MaskModule without refinement does not yield significant improvements. The DepthModule refinement gives a major improvement. The best performance is achieved by combining all the proposed components.



Figure 7: **Comparison on Moving Objects Depth Estimation**: In comparison to other MVS methods, MonoRec is able to predict plausible depths. Furthermore, the depth prediction has less noise and artifacts in static regions of the scene.

strates very different motion and image quality compared to KITTI. The results are shown in Figure 8. The monocular methods struggle to generalize to a new context. The compared MVS methods show more artifacts and cannot predict plausible depths for the moving objects. In contrast our method is able to generalize well to the new scenes for both depth and moving object predictions. Since Oxford RobotCar also provides LiDAR depth data, we further show a quantitative evaluation in the supplementary material.

5. Conclusion

We have presented MonoRec, a deep architecture that estimates accurate dense 3D reconstructions from only a single moving camera. We first propose to use SSIM as the photometric measurement to construct the cost volumes. To deal with dynamic objects, we propose a novel MaskModule which predicts moving object masks from the input cost volumes. With the predicted masks, the proposed DepthModule is able to estimate accurate depths for both static and dynamic objects. Additionally, we propose a novel multi-stage training scheme together with a semisupervised loss formulation for training the depth prediction. All combined, MonoRec is able to outperform the state-of-the-art MVS and monocular depth prediction methods both qualitatively and quantitatively on KITTI and also shows strong generalization capability on Oxford RobotCar and TUM-Mono. We believe that this capacity to recover accurate dense 3D reconstructions from a single moving camera will help to establish the camera as the lead sensor for autonomous systems.



Figure 8: **Oxford RobotCar and TUM-Mono**: All results are obtained by the respective best-performing variant in Table 1. MonoRec shows stronger generalization capability than the monocular methods. Compared to DeepMVS and DeepTAM, MonoRec delivers depth maps with less artifacts and predicts the moving object masks in addition.

Acknowledgement This work was supported by the Munich Center for Machine Learning and by the ERC Advanced Grant SIMULACRON.

References

- M. Bloesch, J. Czarnowski, R. Clark, S. Leutenegger, and A. J. Davison. CodeSLAM - learning a compact, optimisable representation for dense visual SLAM. In *IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), pages 2560–2568, 2018.
- [2] Neill D. F. Campbell, George Vogiatzis, Carlos Hernández, and Roberto Cipolla. Using multiple hypotheses to improve depth-maps for multi-view stereo. In *European Conference* on Computer Vision (ECCV), pages 766–779, 2008.
- [3] Rui Chen, Songfang Han, Jing Xu, and Hao Su. Point-based multi-view stereo network. In *International Conference on Computer Vision (ICCV)*, 2019.
- [4] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 3213–3223, 2016.
- [5] Jan Czarnowski, Tristan Laidlow, Ronald Clark, and Andrew J. Davison. DeepFactors: Real-time probabilistic dense monocular SLAM. *IEEE Robotics and Automation Letters* (*RA-L*), 5(2):721–728, 2020.
- [6] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *International Conference on Computer Vision (ICCV)*, pages 2650–2658, 2015.
- [7] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Neural Information Processing Systems (NIPS)*, 2014.
- [8] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 40(3):611–625, 2018.
- [9] Jakob Engel, Vladyslav Usenko, and Daniel Cremers. A photometrically calibrated benchmark for monocular visual odometry. In *arXiv*, July 2016.
- [10] Alejandro Fontan, Javier Civera, and Rudolph Triebel. Information-driven direct rgb-d odometry. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4929–4937, 2020.
- [11] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), pages 2002–2011, 2018.
- [12] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, pages 1362– 1376, 2010.
- [13] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *International Conference on Computer Vision* (ICCV), 2015.
- [14] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *Inter-*

national Journal of Robotics Research (IJRR), pages 1229–1235, 2013.

- [15] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361. IEEE, 2012.
- [16] Clement Godard, Oisin Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with leftright consistency. In *IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2017.
- [17] Clément Godard, Oisin Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth estimation. In *International Conference on Computer Vision (ICCV)*, pages 3828–3838, 2019.
- [18] Ariel Gordon, Hanhan Li, Rico Jonschkowski, and Anelia Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *International Conference on Computer Vision (ICCV)*, 2019.
- [19] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [20] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3D packing for self-supervised monocular depth estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2485– 2494, 2020.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pages 770–778, 2016.
- [22] Yuxin Hou, Juho Kannala, and Arno Solin. Multi-view stereo by temporal nonparametric fusion. In *International Conference on Computer Vision (ICCV)*, 2019.
- [23] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. DeepMVS: Learning multi-view stereopsis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2821–2830, 2018.
- [24] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. DPSNet: End-to-end deep plane sweep stereo. In *International Conference on Learning Representations (ICLR)*, 2019.
- [25] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In *Neu*ral Information Processing Systems (NIPS), pages 2017– 2025, 2015.
- [26] Mengqi Ji, Jürgen Gall, Haitian Zheng, Yebin Liu, and Lu Fang. SurfaceNet: An end-to-end 3D neural network for multiview stereopsis. In *International Conference on Computer Vision (ICCV)*, pages 2326–2334, 2017.
- [27] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *Neural Information Processing Systems (NIPS)*, page 364–375, 2017.
- [28] Kiriakos N. Kutulakos and Steven M. Seitz. A theory of shape by space carving. In *International Conference on Computer Vision (ICCV)*, 1999.

- [29] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *International Conference on 3D Vision (3DV)*, 2016.
- [30] Maxime Lhuillier and Long Quan. A quasi-dense approach to surface reconstruction from uncalibrated images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (*PAMI*), pages 418–433, 2005.
- [31] 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 Conference on Computer Vision* and Pattern Recognition (CVPR), 2015.
- [32] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4521–4530, 2019.
- [33] Keyang Luo, Tao Guan, Lili Ju, Haipeng Huang, and Yawei Luo. P-MVSNet: Learning patch-wise matching confidence aggregation for multi-view stereo. In *International Conference on Computer Vision (ICCV)*, 2019.
- [34] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. 39(4), 2020.
- [35] Will Maddern, Geoff Pascoe, Chris Linegar, and Paul Newman. 1 Year, 1000km: The Oxford RobotCar Dataset. *International Journal of Robotics Research (IJRR)*, 36(1):3–15, 2017.
- [36] Zak Murez, Tarrence van As, James Bartolozzi, Ayan Sinha, Vijay Badrinarayanan, and Andrew Rabinovich. Atlas: Endto-end 3D scene reconstruction from posed images. In *European Conference on Computer Vision (ECCV)*, 2020.
- [37] Richard A. Newcombe, Steven J. Lovegrove, and Andrew J. Davison. DTAM: Dense tracking and mapping in real-time. In *International Conference on Computer Vision (ICCV)*, 2011.
- [38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In Advances in neural information processing systems, pages 8026–8037, 2019.
- [39] Matia Pizzoli, Christian Forster, and Davide Scaramuzza. REMODE: Probabilistic, monocular dense reconstruction in real time. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [40] Rene Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun. Dense monocular depth estimation in complex dynamic scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4058–4066, 2016.
- [41] Andrea Romanoni and Matteo Matteucci. TAPA-MVS: Textureless-aware PAtchMatch multi-view stereo. In International Conference on Computer Vision (ICCV), 2019.
- [42] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Com-*

puting and Computer Assisted Intervention (MICCAI), pages 234–241. Springer, 2015.

- [43] Chris Russell, Rui Yu, and Lourdes Agapito. Video pop-up: Monocular 3d reconstruction of dynamic scenes. In *European Conference on Computer Vision (ECCV)*, pages 583– 598. Springer, 2014.
- [44] Johannes L Schonberger and Jan-Michael Frahm. Structurefrom-motion revisited. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, 2016.
- [45] Johannes L. Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, pages 501–518, 2016.
- [46] Steven M. Seitz and Charles R. Dyer. Photorealistic scene reconstruction by voxel coloring. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1997.
- [47] Jan Stühmer, Stefan Gumhold, and Daniel Cremers. Realtime dense geometry from a handheld camera. In DAGM Conference on Pattern Recognition, pages 11–20, 2010.
- [48] Keisuke Tateno, Federico Tombari, Iro Laina, and Nassir Navab. CNN-SLAM: Real-time dense monocular SLAM with learned depth prediction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [49] Engin Tola, Christoph Strecha, and Pascal Fua. Efficient large-scale multi-view stereo for ultra high-resolution image sets. *Machine Vision and Applications (MVA)*, pages 903– 920, 2011.
- [50] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant CNNs. In *International Conference on 3D Vision (3DV)*, pages 11– 20. IEEE, 2017.
- [51] Vladyslav Usenko, Nikolaus Demmel, David Schubert, Jörg Stückler, and Daniel Cremers. Visual-inertial mapping with non-linear factor recovery. *IEEE Robotics and Automation Letters (RA-L)*, 5(2):422–429, 2020.
- [52] Chaoyang Wang, Jose Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [53] Kaixuan Wang and Shaojie Shen. MVDepthNet: Real-time multiview depth estimation neural network. In *International Conference on 3D Vision (3DV)*, 2018.
- [54] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [55] Jamie Watson, Oisin Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2021.
- [56] Youze Xue, Jiansheng Chen, Weitao Wan, Yiqing Huang, Cheng Yu, Tianpeng Li, and Jiayu Bao. MVSCRF: Learning multi-view stereo with conditional random fields. In *International Conference on Computer Vision (ICCV)*, 2019.

- [57] Jiayu Yang, Wei Mao, Jose M. Alvarez, and Miaomiao Liu. Cost volume pyramid based depth inference for multi-view stereo. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [58] Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. D3VO: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [59] Nan Yang, Rui Wang, Jörg Stückler, and Daniel Cremers. Deep virtual stereo odometry: Leveraging deep depth prediction for monocular direct sparse odometry. In *European Conference on Computer Vision (ECCV)*, pages 817–833, 2018.
- [60] Yao Yao, Shiwei Li, Siyu Zhu, Hanyu Deng, Tian Fang, and Long Quan. Relative camera refinement for accurate dense reconstruction. In *International Conference on 3D Vision* (3DV), 2017.
- [61] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. MVSNet: Depth inference for unstructured multiview stereo. In *European Conference on Computer Vision* (ECCV), pages 785–801, 2018.
- [62] Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent MVSNet for high-resolution multi-view stereo depth inference. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [63] Zhichao Yin and Jianping Shi. GeoNet: Unsupervised learning of dense depth, optical flow and camera pose. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2018.
- [64] Zehao Yu and Shenghua Gao. Fast-MVSNet: Sparse-todense multi-view stereo with learned propagation and gaussnewton refinement. In *IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2020.
- [65] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian M. Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [66] Huizhong Zhou, Benjamin Ummenhofer, and Thomas Brox. DeepTAM: Deep tracking and mapping. In *European Conference on Computer Vision (ECCV)*, pages 822–838, 2018.
- [67] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.