

CamLessMonoDepth: Monocular Depth Estimation with Unknown Camera Parameters

Sai Shyam Chanduri¹
chanduriss@gmail.com

Zeeshan Khan Suri¹
zshn25@gmail.com

Igor Vozniak²
Igor.Vozniak@dfki.de

Christian Müller²
Christian.Mueller@dfki.de

¹ University of Applied Sciences (HTW)
Saarbrücken, Germany

² Department of Agents and Simulated
Reality (ASR)
German Research Center for Artificial
Intelligence (DFKI GmbH)
Saarbrücken, Germany

Abstract

Perceiving 3D information is of paramount importance in many applications of computer vision. Recent advances in monocular depth estimation have shown that gaining such knowledge from a single camera input is possible by training deep neural networks to predict inverse depth and pose, without the necessity of ground truth data. The majority of such approaches, however, require camera parameters to be fed explicitly during training. As a result, image sequences from wild cannot be used during training. While there exist methods which also predict camera intrinsics, their performance is not on par with novel methods taking camera parameters as input. In this work, we propose a method for implicit estimation of pinhole camera intrinsics along with depth and pose, by learning from monocular image sequences alone. In addition, by utilizing efficient sub-pixel convolutions, we show that high fidelity depth estimates can be obtained. We also embed pixel-wise uncertainty estimation into the framework, to emphasize the possible applicability of this work in practical domain. Finally, we demonstrate the possibility of accurate prediction of depth information without prior knowledge of camera intrinsics, while outperforming the existing state-of-the-art approaches on KITTI benchmark.

1 Introduction

Perceiving accurate depth is a prerequisite for many application domains like robotics and autonomous driving. Traditionally, such systems rely on information fused from depth sensors such as LiDAR, because of their accuracy and robustness. However, such sensors have limited applicability in extreme weather conditions (fog or heavy rain), limited range, higher cost, and complexity [1, 2]. Estimating depth from RGB images can partially mitigate some of these limitations [3]. While supervised learning methods [4, 5, 6, 7, 8] have shown to estimate depth without sensors during inference, they still need ground truth supervision to complete the training. However, acquiring such a large ground truth dataset is

a formidable challenge, moreover, expensive and time-consuming. The utilization of synthetic datasets with "free" ground truth data is a possible solution. However, such datasets are lacking representability compared to natural dataset, owing the difficulty in generating photo-realistic synthetic images [85]. Recent advancements in deep learning helped in alleviating these issues by estimating depth from unlabelled image sequences alone, leading to self-supervised approaches like [20, 21, 60, 64].

Monocular self-supervised approaches offer a more attractive solution than stereo approaches because of the widespread availability of image sequences available for training. In addition, these monocular approaches require no synchronization of cameras. However, they come with their own set of challenges. They operate under a static scene assumption, meaning, the camera is moving, and the scene is static and violation of such assumption predicts 'holes' in depth maps [21]. They also suffer from issues including scaling to metric-depth when ground truth data is not available, brightness changes because of non-Lambertian and reflective surfaces and, occlusions. Over time, many methods with complex network architectures [23, 64], engineered loss functions [21, 49], masking moving objects [21, 36, 64], using ground truth supervision [23] or post-processing [39, 57] for addressing scale ambiguity, etc., have tried to address such challenges. Yet, current monocular approaches significantly lag behind their counterpart supervised approaches like [2, 3, 15].

Another main limitation in many of the earlier approaches [2, 21, 60, 64] is the necessity of precise camera calibration parameters as input, for training accurate depth estimation models. This eliminates the plausible usage of potential data from the wild for training. Previous works [8, 22, 51] have addressed such issues in the past, however, are not on par with the other approaches which use ground-truth calibrated data. In our work, we primarily address this issue by removing the necessity of pre-calibrated data and also focus on refining monocular depth estimation accuracy. The contributions of this work are fourfold: **First**, we learn depth from monocular image sequences even when camera intrinsics are unknown, in a self-supervised fashion, motivated from [22]. We combine it with minimum re-projection loss and auto masking, as proposed by Godard *et al.* [21], in order to deal with occlusions and static pixels. **Second**, we reiterate the importance of using Efficient Sub-Pixel Convolution Networks (ESPCN) (adapted from [11, 48]) for upsampling purposes to obtain sharper depth estimates. Such a method can leverage super-resolution, that can be more accurate in comparison with interpolation approaches [43]. **Third**, we extend our approach to estimate heteroscedastic pixel-wise uncertainty for depth map, capturing by brightness changes due to specular regions or sudden illumination changes, etc. Such information can be leveraged by the agents to take optimal decisions when they are under-confident in predicting depth values in unknown environments, which can avoid fatalities. For example, in an autonomous driving scenario, uncertainty estimation could be beneficial in handling erroneous estimations from the system to prevent accidents [43]. **Finally**, with our exhaustive experiments, we demonstrate that our models, to our best knowledge, outperform the state-of-the-art in monocular depth estimation and closes-in the gap with full-supervised methods on the standard KITTI benchmark [18] even without the need of camera intrinsics as input.

2 Related Work

Monocular Depth Estimation. Monocular depth estimation is an ill-posed problem, as multiple plausible depths could correspond to the same pixel on the image plane. Recent works showed the possibility of accurate depth estimation by analysing patterns in appearances using image sequences alone. For instance, Eigen *et al.* [12] proposed the first deep

learning method using multiscale CNNs which regressed depth output by taking only a single image as input. Many such subsequent supervised approaches like [15, 24, 54] were later proposed, which further extended this formulation. These approaches demand huge amounts of training data, where obtaining ground truth data is both expensive and time-consuming. To mitigate this problem, [58] was proposed which makes the use of synthetic data for generating ground truth. However, it lacks representability compared to the natural training data. To tackle this constraint, [16, 20] were proposed by utilizing the self-supervised strategy that involved learning depth using stereo images for monocular depth inference. In a nutshell, the right images were warped onto the left images using a differentiable sampler as in [26] which enables learning the depth in an end-to-end manner. Later, Zhou *et al.* [64] extended this idea by proposing a strategy to learn the depth along with the pose, to handle training completely with monocular settings only. Furthermore, Godard *et al.* [21] proposed another set of extensions, where the multiscale approach and per-pixel minimum reprojection loss were adopted for the better handling of occlusions. Other self-supervised monocular depth estimation methods like [0, 23, 42, 45, 52, 50] were proposed over the time, which included more robust architectures, additional loss terms and constraints.

Learning from Videos in Wild. Most of the self-supervised learning works [0, 21, 50, 54] require camera intrinsics to learn depth, making it difficult to train on multiple datasets at a time. While Ranftl *et al.* [44] proposed a method to train on multiple datasets at once, their approach to learn depth is supervised, requiring ground-truth depth maps. The lack of camera parameters in generic videos, for instance, a YouTube video, where image sequences are captured from unknown camera setups, limits the usage of such data for training. Traditionally, camera parameters are estimated using techniques which involve calibration targets [40, 47, 51], geometric structures [9, 52] or separate neural networks [5, 14] and are fed as input to the depth estimation models. The main disadvantages of using such approaches include either the necessity of additional data (for calibration), geometric assumptions, or additional complexity and training time. A recent work, proposed by Gordon *et al.* [22] eliminates this need of pre-calibrated camera parameters as input for depth estimation, where camera parameters along with depth, pose and object motion are simultaneously learnt in a fully self-supervised manner. As camera parameters are learnt in a fully self-supervised manner, even image sequences from wild can be used, where the network generalizes better. Later, other works, applicable to image sequences from wild, such as in [8] and [51] were proposed which followed an approach similar to [22] in camera parameters estimation. Our approach is similar to these, but does not need object motion mask as input [22, 51] or online refinement [8], to output accurate depth maps, while assuming a pinhole camera model with minimum or no distortion.

Depth Super-resolution. The problem of image enhancement has been a primary challenge in the fraternity of computer vision. Approaches like [9, 30, 34, 50] have shown to solve this problem while improving upon traditional methods like interpolation. However, such methods demand huge training times and are infeasible when used in combination with applications like depth estimation. For faster and accurate results, Shi *et al.* [48] proposed ESPCN and later [42] has extended it to depth estimation, but requires camera parameters as input. In our work, we exploit ESPCN along with initialization technique from [0] to remove checkerboard artifacts. Consequently, it improves depth estimation accuracy and also, enables faster training and inference compared to other interpolation approaches [48].

Uncertainty. The two types of uncertainties in Bayesian world, namely epistemic and aleatoric, are discussed in recent works [28, 29]. Later Klodt *et al.* [52] proposed an approach to model aleatoric uncertainty which also served for increasing depth estimation accuracy.

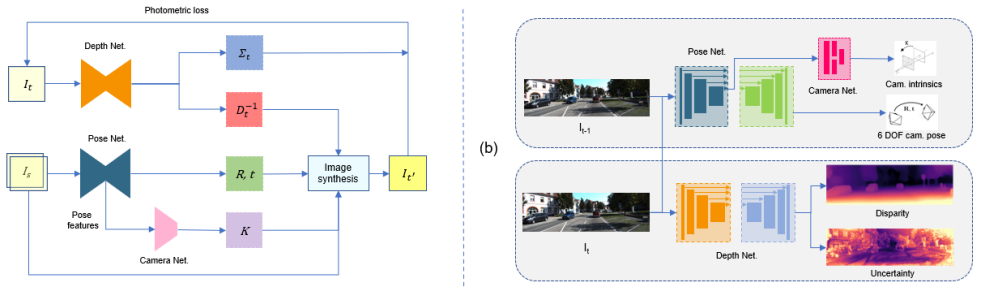


Figure 1: (a) shows the overall pipeline used in our approach. (b) depicts the inference model where the depth model can be used to predict inverse-depth and uncertainty from a single image. The pose and camera models can be used to predict camera pose and intrinsics respectively from two temporal frames.

In this paper, the authors have used a Bayesian framework to model photometric uncertainty as predictive variance. Another approach [27] made use of discrete disparity volume module and to model depth uncertainties. Later, Poggi *et al.* [43] explored various approaches including both empirical and predictive techniques for depth uncertainty estimation in self-supervised methods. Inspired from [49], we propose an approach to model heteroscedastic aleatoric uncertainty as photometric variance by considering noise inherent in the input data.

3 Methodology

Self-supervised depth estimation approaches exempt the necessity of hard ground truth by generating a supervision signal using a moving camera setup and by measuring the overlap between temporal image sequences. This Structure-From-Motion (SFM) problem is modelled by performing a novel image synthesis of the target frame from image source point of view. However, this objective is ill-posed since knowledge of accurate pose and camera intrinsics can reconstruct this novel view even with incorrect depth estimations [21]. To address such ambiguity, appearance matching and smoothness losses were proposed. However, it demands the knowledge of pre-calibrated camera intrinsics for novel view synthesis.

Motivated from earlier approaches [21, 22], we propose a framework to simultaneously learn depth, camera pose, uncertainty and camera intrinsics from input video sequences, as seen in Figure 1(a). The depth network predicts a disparity map (inverse depth) D_t^{-1} of the target image I_t and its corresponding uncertainty map Σ_t . Simultaneously, the pose network takes any two consecutive frames I_s (either I_{t+1} or I_{t-1}) and I_t , outputs a 6 Degrees-of-Freedom (DOF) rigid transformation $T_{t \rightarrow s} \in \text{SE}(3)$ from the target image plane to the source image plane, containing rotational $R \in \text{SO}(3)$ and translational $t \in \mathbb{R}^3$ information. The camera network takes the pose encoder features as input and outputs the camera intrinsics matrix K , containing principal offset and focal length information. After obtaining these predictions, we reconstruct I_t from the source image's point of view (POV). This novel view synthesis involves lifting the target image into 3D using the predicted depth map D_t and the intrinsics matrix's inverse K^{-1} . Then, using the predicted transformation matrix $T_{t \rightarrow s}$, and K , we project the 3D scene onto the source's POV, thus warping the source image into the target image [21, 64]. We call this warped target image $I_{t'}$. Such a transformation from the target image homogeneous coordinates p_t to the source ones $p_s \in \mathbb{R}^3$, is summarized by the relation $p_s \stackrel{\dagger}{=} \langle KRK^{-1}D_t(p_t)p_t + Kt \rangle$ [64], where $\langle \rangle$ denotes the sampling operator, forms the backbone of this approach.

Relying on a static scene assumption, the warped target image $I_{t'}$ is compared against the original target image I_t , using a photometric error metric $\mathcal{L}_p(I_t, I_{t'})$, which is minimized. Following previous approaches [21, 64], we use a convex combination of L1 loss and a structural similarity(SSIM) loss [65] for calculating the photometric error, with $\alpha = 0.85$. In order to handle occlusions, we follow Godard *et al.*'s [21] use of per-pixel minimum reprojection which compares views generated from multiple source frames to the target frame and considers the minimum of the per-pixel photometric error, as in Equation 1. This computation prevents high error values when correspondences are good [21].

$$\mathcal{L}_p(I_t, I_s) = \min_{I_s} \mathcal{L}_p(I_t, I_{t'}) = \min_{I_s} \left(\alpha \frac{1 - \text{SSIM}(I_t, I_{t'})}{2} + (1 - \alpha) \|I_t - I_{t'}\|_1 \right) \quad (1)$$

We learn a heteroscedastic Bayesian uncertainty map Σ_t , for input I_t , $\sigma_{ij} \in [0, 1] \forall \sigma_{ij} \in \Sigma_t$, capturing the photometric variance caused by specular objects or by sudden illumination change, etc., following from Kendall *et al.*'s [49] residual weighting, as shown in Equation 2. Because of such modelling, the pixels corresponding to such brightness changes are down weighted [69] and as a result, the uncertainty values are higher at regions with high photometric reprojection error.

$$\mathcal{L}_p(I_t, I_s)^* = \frac{\min_{I_s} \mathcal{L}_p(I_t, I_{t'})}{2\Sigma_t^2} + \frac{1}{2} \log \Sigma_t + 1.5 \quad (2)$$

Real world scenes do not always follow a static scene assumption, which is necessary for self-supervised monocular depth estimation. The performance of these methods suffers greatly in presence of static camera or object motion in same direction as camera which can manifest infinite depth results [21]. To handle such scenarios where the relative motion between camera and objects is zero, we make use of binary mask $\mu \in \{0, 1\}$ from [21] to filter out static pixels shown in Equation 3. Here $[\]$ represents the Iverson bracket. Such static pixels are removed by identifying pixels with higher warped reprojection loss between the warped source frame $I_{t'}$ and target frame I_t than the unwarped reprojection loss between the source and target frames.

$$\mu = \left[\min_{I_s} \mathcal{L}_p(I_t, I_s) > \min_{I_s} \mathcal{L}_p(I_t, I_{t'}) \right] \quad (3)$$

In order to encourage the smoothness of disparity estimations, especially in textureless and low-image gradient areas, and for regularizing inverse depth to prevent divergent values, we add a combination of first and second order smoothness losses \mathcal{L}_s . Equation 4 represents such a disparity smoothness term, weighted by the image's gradients, in order to preserve edges. The first-order gradient term (from) $\nabla_1 \equiv \partial_x + \partial_y$, is an L1 penalty on disparity gradients used to account for the depth discontinuities which often occur at strong image gradients [49]. Adding a second order gradient $\nabla_2 \equiv \partial_{xx} + \partial_{xy} + \partial_{yx} + \partial_{yy}$, in addition, encourages better smoothness of these gradients with larger convergence radii, leading to better optimization [49]. To discourage shrinking of the estimated depth [21], mean normalized inverse depth $\hat{D}_t^* = D_t^{-1} / \bar{D}_t$ is considered, where D_t^{-1} represents inverse depth or disparity and \bar{D}_t represents the mean disparity.

$$\mathcal{L}_s = |\nabla_1 \hat{D}_t^*| e^{-0.5|\nabla_1 I_t|} + |\nabla_2 \hat{D}_t^*| e^{-0.5|\nabla_2 I_t|} \quad (4)$$

Final Training Loss. The final SFM objective now, is the combination of the net per-pixel minimum photometric loss with uncertainty (Equation 2) and the disparity smoothness loss (Equation 4) given by $\mathcal{L}_t = \mu \mathcal{L}_p^* + \lambda \mathcal{L}_s$. Here, λ is the smoothness regularizer, to be chosen as a hyperparameter. This total loss \mathcal{L}_t is averaged over the total number of scales used in depth decoder, which is 4 in our work.

Network details. The depth and pose networks follows U-net type encoder-decoder architecture with skip connections in between the encoder and the decoder which facilitates better learning of deep abstract features along with spatial information. We use ResNext-50 [56] **depth encoder** with a cardinality of 32 and base width of 4 and ResNet-50 **pose encoder**, unless stated otherwise. Following earlier works like [21, 24, 27, 53] in depth estimation, we also employ pre-trained ImageNet [46] weights for initialization for the encoder. However, we use weights pre-trained in a semi-weakly supervised fashion as proposed in [58] for both depth and pose networks unlike earlier works [24, 21, 23], which has improved depth estimation accuracy for us. **Depth decoder** uses a multiscale architecture to output inverse-depth and uncertainty at 4 scales similar to [21] with two primary modifications: (i) nearest-neighbour interpolation replaced with ESPCN for upsampling following ideas from [48] and [1] and, (ii) modification of final layer in the decoder to output depth uncertainty along with disparity. **Pose decoder** follows architecture from [21]. **The camera network** is inspired from [22] but the architecture is slightly different. More network details are mentioned in supplementary section. The features obtained from the pose encoder are passed to this camera network, followed by squeeze operation to reduce the number of channels to 256. Two independent 3x3 convolution layers stem from this squeeze layer to estimate the normalized focal lengths f_x, f_y and principal offsets c_x, c_y in horizontal and vertical axes, normalized by the input image’s width and height respectively. These are concatenated to output the camera intrinsics matrix $K = \begin{pmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{pmatrix}$. The softplus activation function $f(x) = \log(1 + \exp(x))$ is used to avert negative values for focal lengths.

Efficient Sub-pixel convolutions. Since we use an encoder-decoder architecture for depth, it is necessary to upsample the outputs in decoder layers. Earlier approaches like [21, 53, 60, 62] rely on a nearest-neighbour interpolation for faster inference. Such techniques could compromise the output at object boundaries, as they combine the values from foreground and background. Moreover, the filters are not learnable, thus, limiting the upsampling effect. Transpose convolutions, on the contrary, are learnable, but suffer from checkerboard artifacts. To compensate for these artifacts, usually they are initialized with interpolation outputs, which come at the cost of training time. Hence, we make use of ESPCN [48] to perform convolutional learning at low resolution. We then perform a pixel shuffle operation to upsample only on the final step. To alleviate the checkerboard artifacts caused by random initialization of these filters, we make use of Initialization to Convolution Nearest-neighbour Resize (ICNR) from [1] which provides an initialization similar to nearest-neighbour. Such operation results in faster convergence and leads to a better minimum compared to transpose convolutions and other interpolation approaches [48].

4 Experiments

Setup. Our PyTorch [41] models were trained on 4 Nvidia GeForce RTX 2080 Ti GPUs, in a distributed setting, using the Adam optimizer [31], with a learning rate of 10^{-4} for the first 15 epochs and 10^{-5} for the last 10. The batch size is set to 12 and the default resolution is 640x192 (width x height), unless stated otherwise. The models are trained with data augmentation with a 50% chance of horizontal flips, random brightness (± 0.2), random

contrast (± 0.2), hue jitters (± 0.1) and, saturation (± 0.2) variations. The input trio of frames at times $t - 1$, t , and $t + 1$, are all applied with the same augmentation settings, but these augmentations are not used for computing the photometric loss. The max. and min. depths are set to 0.1 and 100 respectively. The remaining training parameters are set as in [24].

4.1 KITTI Dataset

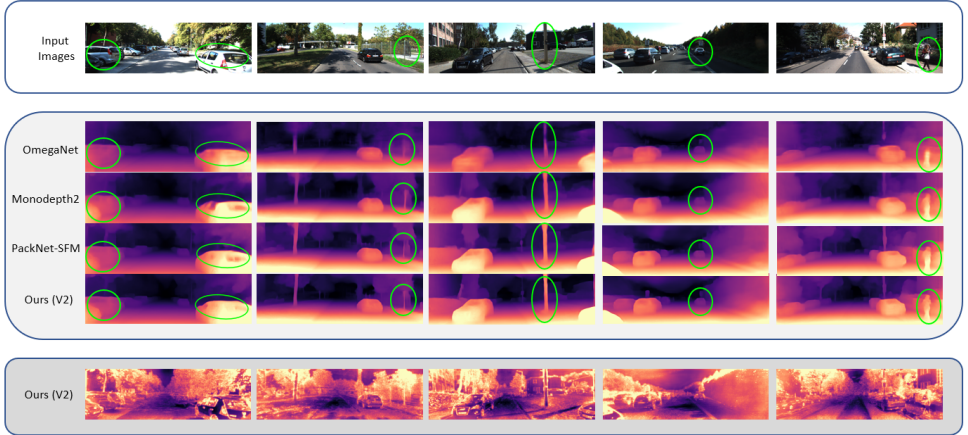


Figure 2: The input images, their corresponding disparity maps, and uncertainty maps (bottom) are shown with blocks. Our depth results are compared with MonoDepth2 [24]), OmegaNet [61] and PackNet-SFM [23]. For disparity maps, brighter pixels (orange) indicate nearer regions while darker indicate for farther. It can be observed that our approach shows high-quality depth maps especially at object boundaries. For uncertainty maps, brighter regions (orange) indicate pixels with higher uncertainty. Such uncertainty is more pronounced along object boundaries, when further from the centre of camera, also where brightness change occur, e.g. at reflective surfaces (a).

Method	Intr	Lower is better				Higher is better		
		Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
SFMlearner [62]	✗	0.208	1.768	6.958	0.283	0.678	0.885	0.957
MonoDepth2 [24]	✗	0.115	0.903	4.863	0.193	0.877	0.959	0.981
PackNet-SFM [23]	✗	0.111	0.785	4.601	0.189	0.878	0.960	0.982
HR-Depth [63]	✗	0.109	0.792	4.632	0.185	0.884	0.962	0.983
Johnston <i>et al.</i> [24]	✗	0.106	0.861	4.699	0.185	0.889	0.962	0.982
GLNet (-ref) [9]	✓	0.135	1.070	5.230	0.210	0.841	0.948	0.980
Gordon <i>et al.</i> [24]	✓	0.128	0.959	5.230	0.212	0.845	0.947	0.976
OmegaNet [61]	✓	0.120	0.792	4.750	0.191	0.856	0.958	0.984
Ours without uncert. (V1)	✓	0.105	0.768	4.539	0.182	0.890	0.964	0.983
Ours with uncert. (V2)	✓	0.106	0.750	4.482	0.182	0.891	0.964	0.983

Table 1: Quantitative results on KITTI test data with Eigen split. Here, "Intr" indicates if camera intrinsics are learned or not during training, "-ref" meaning without refinement. The results are taken from their corresponding papers. Best results are highlighted in bold.

The KITTI benchmark has become a de facto standard for depth evaluation in recent times [23]. Hence, for evaluation, we make use of KITTI 2015 dataset [24] which con-

sists of 200 different scenes of driving data captured using RGB cameras along with sparse ground truth depth maps captured by a Velodyne LiDAR sensor. As the pre-processing step, we follow the previous works [64] to remove static frames. This resulted in 39810 monocular raw images for training and 4424 for validation. In all the evaluation experiments, we make use of seven different metrics as proposed by Eigen *et al.* [43], which are now commonly used and accepted evaluation indicators for comparison among various depth estimation approaches [63]. A depth cap of 80 meters is used in all our experiments.

Evaluation. In this section, we compare the effectiveness of the proposed method with SOTA using KITTI test dataset and Eigen split [44]. Eigen split contains 697 images with reprojected LIDAR points, which are used for evaluation. For depth evaluation, per-image median ground truth scaling (as proposed in [44]) is utilized, to handle the unavailability of scale information in self-supervised monocular approaches [24, 64]). No further post-processing steps were involved. We compare the results of our two variants, one without uncertainty estimation (model V1) and the other with uncertainty estimation (model V2) for evaluation on KITTI benchmark. As shown in Table 1, both our models achieved significant gains in both uncalibrated and calibrated scenarios in comparison to their corresponding baseline methods - OmegaNet [61] and MonoDepth2 [24] respectively. In an uncalibrated scenario, where ground truth camera intrinsics are not known, our camera network predicts such values taking in a pair of images each time. Using such method, along with the proposed considerations, produced results which outperform the SOTA not only in uncalibrated but also in calibrated settings. In addition, a vivid improvement in depth map quality over baseline methods can be observed from qualitative results shown in Figure 2. This shows the effectiveness of the proposed method, especially at object boundaries, primarily due to the proposed ESPCN layers. Notably, we achieved such competing metric scores using ResNext-50 [66] encoder, on contrary to high-end architectures used in [23] and [27].

Method	Resolution	Lower is better				Higher is better		
		Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
OmegaNet [61]	640x192	0.120	0.792	4.750	0.191	0.856	0.958	0.984
MonoDepth2 [24]	640x192	0.115	0.903	4.863	0.193	0.877	0.959	0.981
PackNet-SfM [62]	640x192	0.111	0.785	4.601	0.189	0.878	0.960	0.982
HR Depth [65]	640x192	0.109	0.792	4.632	0.185	0.884	0.962	0.983
Ours	640x192	0.106	0.750	4.482	0.182	0.891	0.964	0.983
OmegaNet [61]	1024x320	0.118	0.748	4.608	0.186	0.865	0.961	0.985
MonoDepth2 [24]	1024x320	0.115	0.882	4.701	0.190	0.879	0.961	0.982
PackNet-SfM [62]	1280 x 384	0.107	0.802	4.538	0.186	0.889	0.962	0.981
HR Depth [65]	1024x320	0.106	0.755	4.472	0.181	0.892	0.966	0.984
HR Depth [65]	1280 x 384	0.104	0.727	4.410	0.179	0.894	0.966	0.984
Ours	1024x320	0.102	0.723	4.374	0.178	0.898	0.966	0.983

Table 2: KITTI benchmark comparison with our baseline methods at different resolutions

Input resolution. In this experimental study, we have considered two common groups with image resolutions of 640x192 and 1024x320 pixels, respectively. In accordance to the reported results in Table 2, the proposed method (model V2) outperforms baselines ([24], [61]) and other novel methods, at lower and higher resolutions. While using a higher input image resolution has not improved results significantly in the baseline methods, it has showed positive effect with our approach. Moreover, this improvement due to a change in resolution is highest among the group. This manifests the effectiveness of efficient sub-pixel convolutions at higher resolutions. Even when regressed to an inferior ResNet-18 pose encoder

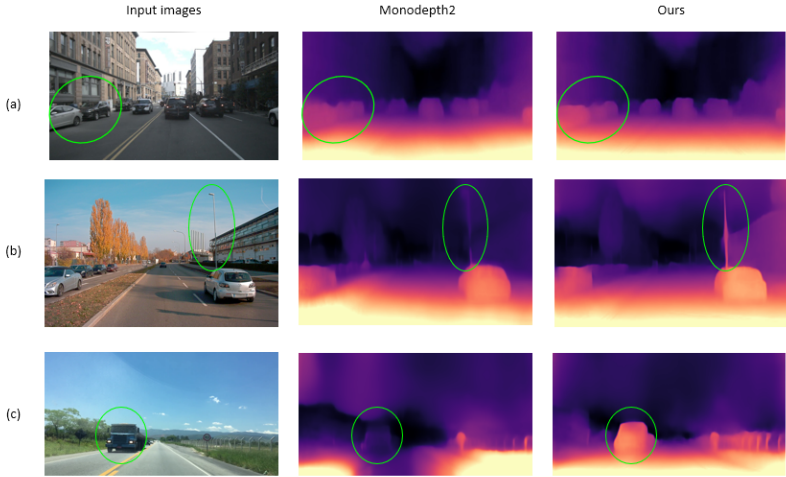


Figure 3: Unseen images from (a) Nuscenes [1], (b) Audi a2d2 [19] datasets respectively, and (c) a random video collected from Pexels. Better boundary separation (a), sharp results on thin objects (b) and accurate depth maps (c) suggest better applicability and generalizability of our model, trained just on KITTI dataset.

Method	Intr	SR	Uncert	Lower is better				Higher is better		
				Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
MD2 (R18) [21]	✗	✗	✗	0.115	0.903	4.863	0.193	0.877	0.959	0.981
MD2 (R50) [21]	✗	✗	✗	0.110	0.831	4.642	0.187	0.883	0.962	0.982
Ours	✗	✗	✗	0.111	0.761	4.743	0.188	0.876	0.960	0.983
Ours	✓	✗	✗	0.110	0.816	4.662	0.187	0.887	0.962	0.982
Ours	✗	✓	✗	0.106	0.792	4.593	0.184	0.890	0.963	0.982
Ours	✓	✓	✗	0.105	0.768	4.539	0.182	0.890	0.964	0.983
Ours	✓	✓	✓	0.106	0.750	4.482	0.182	0.891	0.964	0.983

Table 3: Ablation studies on KITTI benchmark with baseline referring to [21], where “Intr” indicates training with intrinsics network. “SR” indicates the usage of super resolution, and “Uncert” stands for training with uncertainty.

architecture at higher resolution owing to computational limitations, the results achieved along all performance metrics are the best among all our models.

KITTI ablation experiments. To evaluate the significance of each of the components proposed in our approach, we have performed ablation studies on KITTI test data with Eigen split (Table 3). We evaluated depth by varying three main components in this study: (i) with intrinsics network implying that external feed of pre-calibrated camera intrinsic parameters is unnecessary, (ii) using efficient sub-pixel convolutions and, (iii) with uncertainty estimation. Our approach with the networks proposed and training camera intrinsics has itself improved results compared to baseline (here [21]). It can be observed that our model trained without any of our components is not better than MonoDepth2, suggesting that these 3 contributions lead to improvement. The main improvement is associated with the usage of efficient sub-pixel convolutions in the depth decoder, resulting in sharper depth maps at higher resolutions. Crucially, this improvement is even more valuable in terms of achieved performance when trained camera network. This finding is a key result for proving the effectiveness of our proposed method to leverage its usage for potentially unlimited datasets including videos from internet. Adding the heteroscedastic aleatoric uncertainty estimation, has not improved

depth estimation significantly, where the results remained roughly the same.

Generalization capability. We perform qualitative evaluation using our model trained only on KITTI dataset, while testing on unseen data from unseen videos similar to [8]. However, we make use of two additional datasets - NuScenes [6] and Audi a2d2 [19], along with a random road scene video from internet. Qualitative results of this experiment are demonstrated in Figure 3. Our model results are particularly exceptional along the object boundaries and in identifying thin objects compared to our baseline method [24]. This evidence supports our claim that our method is generalizable to learn structures and scenes collected using different cameras, and also without the knowledge of camera intrinsics.

4.2 Videos from Internet



Figure 4: Test ^(a) results of our model, trained on ^(b) multiple image sequences including KITTI and random internet videos. ^(c)

For demonstrating that our method can be used for videos even without ground-truth intrinsics, we have gathered 57 stock videos from the internet, predominantly covering road scenes which involve minimal distortion. The URLs of these videos will be made public. After pre-processing, it resulted in 7623 and 1373 images for training and test respectively. This dataset is particularly challenging as it involves different camera setups, also with varying camera heights, obstacles present throughout some videos (for *e.g.* car hood), highly varying image resolutions, and even weather conditions (day, rain, snow, night *etc.*) as they are from the wild. For training, these images are resized to have a width of 640 and then, centre cropped to 640x192. This step assures that the input image is not stretched, although some vertical field of view is lost. Images from KITTI 2015 dataset [24], without intrinsics, involving similar pre-processing steps are used together for training along with this internet data to handle the variability and unavailability of large training dataset from the internet alone. Qualitative results, are demonstrated in Figure 4, justify that the proposed approach can be used even when camera parameters are not known.

5 Conclusion

We propose a self-supervised monocular approach to learn depth for pinhole camera with minimum distortion, even without explicitly given intrinsics. This makes it possible to utilize videos from wild for training, which can be leveraged to obtain potentially unlimited data. Besides, we demonstrate that by incorporating ESPCN instead of interpolation in the depth decoder, robust and sharper depth maps can be obtained. In addition, we include an approach to estimate pixel-wise depth uncertainties, which could play a crucial role in robotics and autonomous driving tasks in tacking optimal decisions. Our contributions lead to significant improvement on KITTI benchmark even when ground truth camera intrinsics data is not available. Furthermore, through our experiments, we show that our best results are obtained at a higher input resolution of 1024x320 and to our knowledge, this result significantly outperforms the other state-of-the-art self-supervised monocular depth estimation methods on KITTI benchmark among methods which neither use ground truth data nor online refinement techniques for training.

Acknowledgements. This research has been supported by the computer vision team in project SpuMo E2D (funding reference number 03EFLSL015) under HTW Saar, and in part by the autonomous driving team (Agents and Simulated Reality department) of German Research Center for Artificial Intelligence (DFKI GmbH, Saarbrücken).

References

- [1] Andrew Aitken, Christian Ledig, Lucas Theis, Jose Caballero, Zehan Wang, and Wenzhe Shi. Checkerboard artifact free sub-pixel convolution: A note on sub-pixel convolution, resize convolution and convolution resize. *arXiv preprint arXiv:1707.02937*, 2017.
- [2] Ibraheem Alhashim and Peter Wonka. High quality monocular depth estimation via transfer learning. *arXiv preprint arXiv:1812.11941*, 2018.
- [3] Lorenzo Andraghetti, Panteleimon Myriokefalitakis, Pier Luigi Dovesi, Belen Luque, Matteo Poggi, Alessandro Pieropan, and Stefano Mattoccia. Enhancing self-supervised monocular depth estimation with traditional visual odometry. In *2019 International Conference on 3D Vision (3DV)*, pages 424–433. IEEE, 2019.
- [4] João Pedro Barreto and Helder Araujo. Geometric properties of central catadioptric line images and their application in calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1327–1333, 2005.
- [5] Oleksandr Bogdan, Viktor Eckstein, Francois Rameau, and Jean-Charles Bazin. Deepcalib: a deep learning approach for automatic intrinsic calibration of wide field-of-view cameras. In *Proceedings of the 15th ACM SIGGRAPH European Conference on Visual Media Production*, pages 1–10, 2018.
- [6] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [7] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8001–8008, 2019.
- [8] Yuhua Chen, Cordelia Schmid, and Cristian Sminchisescu. Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In *Proceedings of the IEEE international conference on computer vision*, pages 7063–7072, 2019.
- [9] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European conference on computer vision*, pages 184–199. Springer, 2014.
- [10] Gregory Dudek and Michael Jenkin. *Computational principles of mobile robotics*. Cambridge university press, 2010.

- [11] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE international conference on computer vision*, pages 2650–2658, 2015.
- [12] David Eigen, Christian Puhersch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27:2366–2374, 2014.
- [13] David Eigen, Christian Puhersch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *arXiv preprint arXiv:1406.2283*, 2014.
- [14] Jose M. Facil, Benjamin Ummenhofer, Huizhong Zhou, Luis Montesano, Thomas Brox, and Javier Civera. CAM-ConvS: Camera-Aware Multi-Scale Convolutions for Single-View Depth. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [15] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2002–2011, 2018.
- [16] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European conference on computer vision*, pages 740–756. Springer, 2016.
- [17] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361. IEEE, 2012.
- [18] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11): 1231–1237, 2013.
- [19] Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi, Xavier Ricou, Rupesh Durgesh, Andrew S Chung, Lorenz Hauswald, Viet Hoang Pham, Maximilian Mühlegg, Sebastian Dorn, et al. A2d2: Audi autonomous driving dataset. *arXiv preprint arXiv:2004.06320*, 2020.
- [20] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 270–279, 2017.
- [21] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3828–3838, 2019.
- [22] Ariel Gordon, Hanhan Li, Rico Jonschkowski, and Anelia Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8977–8986, 2019.

- [23] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2485–2494, 2020.
- [24] Xiaoyang Guo, Hongsheng Li, Shuai Yi, Jimmy Ren, and Xiaogang Wang. Learning monocular depth by distilling cross-domain stereo networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 484–500, 2018.
- [25] Philipp Heise, Sebastian Klose, Brian Jensen, and Alois Knoll. Pm-huber: Patchmatch with huber regularization for stereo matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2360–2367, 2013.
- [26] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. *arXiv preprint arXiv:1506.02025*, 2015.
- [27] Adrian Johnston and Gustavo Carneiro. Self-supervised monocular trained depth estimation using self-attention and discrete disparity volume. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4756–4765, 2020.
- [28] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pages 5574–5584, 2017.
- [29] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.
- [30] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1637–1645, 2016.
- [31] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [32] Maria Klodt and Andrea Vedaldi. Supervising the new with the old: learning sfm from sfm. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 698–713, 2018.
- [33] Yevhen Kuznietsov, Jorg Stuckler, and Bastian Leibe. Semi-supervised deep learning for monocular depth map prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6647–6655, 2017.
- [34] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [35] Adrian Lopez-Rodriguez and Krystian Mikolajczyk. Desc: Domain adaptation for depth estimation via semantic consistency. *arXiv preprint arXiv: 2009.01579*, 2020.

- [36] Yue Luo, Jimmy Ren, Mude Lin, Jiahao Pang, Wenxiu Sun, Hongsheng Li, and Liang Lin. Single view stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 155–163, 2018.
- [37] Xiaoyang Lyu, Liang Liu, Mengmeng Wang, Xin Kong, Lina Liu, Yong Liu, Xinxin Chen, and Yi Yuan. Hr-depth: High resolution self-supervised monocular depth estimation. *arXiv preprint arXiv:2012.07356*, 2020.
- [38] Nikolaus Mayer, Eddy Ilg, Philipp Fischer, Caner Hazirbas, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. What makes good synthetic training data for learning disparity and optical flow estimation? *International Journal of Computer Vision*, 126(9):942–960, 2018.
- [39] Robert McCraith, Lukas Neumann, and Andrea Vedaldi. Calibrating self-supervised monocular depth estimation. *arXiv preprint arXiv:2009.07714*, 2020.
- [40] Christopher Mei and Patrick Rives. Single view point omnidirectional camera calibration from planar grids. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 3945–3950. IEEE, 2007.
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [42] Sudeep Pillai, Rareş Ambruş, and Adrien Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9250–9256. IEEE, 2019.
- [43] Matteo Poggi, Filippo Aleotti, Fabio Tosi, and Stefano Mattoccia. On the uncertainty of self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3227–3237, 2020.
- [44] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.
- [45] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J Black. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 12240–12249, 2019.
- [46] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

- [47] Shishir Shah and JK Aggarwal. A simple calibration procedure for fish-eye (high distortion) lens camera. In *Proceedings of the 1994 IEEE international Conference on Robotics and Automation*, pages 3422–3427. IEEE, 1994.
- [48] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016.
- [49] Chang Shu, Kun Yu, Zhixiang Duan, and Kuiyuan Yang. Feature-metric loss for self-supervised learning of depth and egomotion. In *European Conference on Computer Vision*, pages 572–588. Springer, 2020.
- [50] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [51] Fabio Tosi, Filippo Aleotti, Pierluigi Zama Ramirez, Matteo Poggi, Samuele Salti, Luigi Di Stefano, and Stefano Mattoccia. Distilled semantics for comprehensive scene understanding from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [52] Sudheendra Vijayanarasimhan, Susanna Ricco, Cordelia Schmid, Rahul Sukthankar, and Katerina Fragkiadaki. Sfm-net: Learning of structure and motion from video. *arXiv preprint arXiv:1704.07804*, 2017.
- [53] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2022–2030, 2018.
- [54] Xiaolong Wang, David Fouhey, and Abhinav Gupta. Designing deep networks for surface normal estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 539–547, 2015.
- [55] 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.
- [56] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [57] F. Xue, G. Zhuo, Z. Huang, W. Fu, Z. Wu, and M. H. Ang. Toward hierarchical self-supervised monocular absolute depth estimation for autonomous driving applications. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2330–2337, 2020. doi: 10.1109/IROS45743.2020.9340802.
- [58] I Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification. *arXiv preprint arXiv:1905.00546*, 2019.

- [59] Nan Yang, Lukas von Stumberg, Rui Wang, and Daniel Cremers. D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1281–1292, 2020.
- [60] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1983–1992, 2018.
- [61] Xianghua Ying and Hongbin Zha. Identical projective geometric properties of central catadioptric line images and sphere images with applications to calibration. *International Journal of Computer Vision*, 78(1):89–105, 2008.
- [62] Mi Zhang, Jian Yao, Menghan Xia, Kai Li, Yi Zhang, and Yaping Liu. Line-based multi-label energy optimization for fisheye image rectification and calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4137–4145, 2015.
- [63] ChaoQiang Zhao, QiYu Sun, ChongZhen Zhang, Yang Tang, and Feng Qian. Monocular depth estimation based on deep learning: An overview. *Science China Technological Sciences*, pages 1–16, 2020.
- [64] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1851–1858, 2017.