ProDepth: Boosting Self-Supervised Multi-Frame Monocular Depth with Probabilistic Fusion

Sungmin Woo*, Wonjoon Lee*, Woo Jin Kim, Dogyoon Lee, and Sangyoun Lee

Yonsei University https://sungmin-woo.github.io/prodepth/

Abstract. Self-supervised multi-frame monocular depth estimation relies on the geometric consistency between successive frames under the assumption of a static scene. However, the presence of moving objects in dynamic scenes introduces inevitable inconsistencies, causing misaligned multi-frame feature matching and misleading self-supervision during training. In this paper, we propose a novel framework called ProDepth, which effectively addresses the mismatch problem caused by dynamic objects using a probabilistic approach. We initially deduce the uncertainty associated with static scene assumption by adopting an auxiliary decoder. This decoder analyzes inconsistencies embedded in the cost volume, inferring the probability of areas being dynamic. We then directly rectify the erroneous cost volume for dynamic areas through a Probabilistic Cost Volume Modulation (PCVM) module. Specifically, we derive probability distributions of depth candidates from both singleframe and multi-frame cues, modulating the cost volume by adaptively fusing those distributions based on the inferred uncertainty. Additionally, we present a self-supervision loss reweighting strategy that not only masks out incorrect supervision with high uncertainty but also mitigates the risks in remaining possible dynamic areas in accordance with the probability. Our proposed method excels over state-of-the-art approaches in all metrics on both Cityscapes and KITTI datasets, and demonstrates superior generalization ability on the Waymo Open dataset.

Keywords: Multi-frame monocular depth estimation \cdot Self-supervised learning \cdot Probabilistic modeling

1 Introduction

Accurate depth information is essential across various domains, including autonomous driving, robotics, and augmented reality. The deployment of precise 3D sensors (e.g., structured light or LiDAR) is often hindered by their high costs, leading to the development of depth estimation solely from RGB images. Notably, self-supervised monocular depth estimation from single or multiple frames is gaining traction, removing the need for ground-truth data from costly sensors.

^{*} co-first authors



Fig. 1: Our ProDepth performs uncertainty-aware adaptive fusion of the probability distributions from both single-frame and multi-frame cues. The fused distribution follows the distribution of single-frame cues for a dynamic pixel, while adhering to the distribution of multi-frame cues for a static pixel. Error maps in the second column depict large depth errors in green and small in blue.

Early self-supervised depth estimation methods [2, 5, 11, 35, 43] take a single target image to infer depth by analyzing visual patterns including texture, shading, and edges. The adjacent images are only incorporated for self-supervision at the training-level, which is achieved by minimizing a photometric reprojection error [43] between frames as a novel-view synthesis problem. However, due to the constraints of limited information, their performance falls short of achieving satisfactory results. Recently, multi-frame based approaches [8, 15, 38] have emerged to leverage temporally adjacent frames as valuable geometric cues for depth estimation. These methods perform multi-frame feature matching within the cost volume under the assumption of a static scene, assessing the probabilities of various depth candidates for each pixel based on geometric consistency between frames. Despite their overall high performance, these approaches exhibit significant errors in dynamic areas. The inconsistent geometric locations of moving objects lead to misaligned feature matching, resulting in an incorrect depth probability distribution.

To address the mismatch problem in the cost volume, several works [8,15,38] leverage single-frame depth to compensate for errors in dynamic areas in multi-frame depth. The underlying insight [25,39] is that multi-frame based estimation tends to yield more accurate predictions in static areas, whereas single-frame based estimation without the cost volume avoids misaligned feature matching, thereby better handling moving objects (Fig. 1). The representative approach [15,38] is to supervise dynamic areas of multi-frame depth with single-frame depth by using an additional training loss term, aiming to enforce correct depth estimation despite an incorrect cost volume. However, this loss-level solution cannot entirely prevent errors in the cost volume from affecting the final prediction, as the fundamental issue of an incorrect multi-frame matching cost

distribution persists. Recently, DynamicDepth [8] indirectly addresses the mismatch issue in the cost volume by adjusting the locations of dynamic objects in input images to be static using single-frame depth. However, this process requires pre-computed segmentation masks to identify objects, and the masks involve needless static objects since segmentation does not account for their movements. While these approaches have made progress in handling dynamic areas, their limitations highlight the need for further exploration. Our key observation is that accurately identifying dynamic objects remains a significant challenge, and the direct refinement of incorrect matching costs in the cost volume has yet to be thoroughly explored.

In this paper, we introduce ProDepth, a novel framework that makes three major contributions to address the inconsistency issue caused by dynamic objects. First, rather than relying on additional semantic information, we discern an uncertainty map (*i.e.*, the probability that each pixel is not static) using an auxiliary depth decoder. This decoder deliberately predicts corrupted depth based on erroneous cost volume, enabling to infer object-level uncertainty from the extent of corruption. Second, we present a Probabilistic Cost Volume Modulation (PCVM) module, which directly rectifies the erroneous matching costs of the cost volume through uncertainty-aware adaptive fusion of single- and multiframe cues. As illustrated in Fig. 1, the depth probability distribution adaptively follows either the distribution of single-frame cues or that of multi-frame cues based on the inferred uncertainty. Finally, to further mitigate the issues associated with incorrect self-supervision of reprojection loss in dynamic areas during training, we devise a loss reweighting strategy. This strategy entails adjusting the computed reprojection loss according to uncertainty, thereby reducing incorrect supervision in possible dynamic areas. In summary, we present the following noteworthy contributions:

- We devise an auxiliary depth decoder, which facilitates the identification of moving objects as a probabilistic representation, *i.e.* uncertainty, without using of a pretrained off-the-shelf segmentation network.
- We propose PCVM, a novel approach addressing the mismatch problem in the cost volume by directly rectifying the corrupted matching cost distribution through the probabilistic fusion of single-frame and multi-frame cues.
- We introduce a self-supervision loss reweighting strategy to counteract incorrect supervision in potential dynamic areas, distinct from conventional binary masking methods.
- Our approach achieves state-of-the-art results on Cityscapes and KITTI datasets, and also demonstrates superior generalization ability on the Waymo Open dataset.

2 Related Work

2.1 Self-Supervised Monocular Depth Estimation

Conventional single-frame based methods use a single image for estimation, with temporally adjacent images employed solely for self-supervision during

training. A self-supervised framework [43] is proposed to compute photometric consistency between monocular frames, facilitating joint training of a singleframe depth estimation network and a multi-frame camera pose estimation network. Subsequent advancements are achieved in camera geometry modeling [13,33], network architectures [14,42], reprojection loss [33,41] and the handling of depth errors in moving objects [2, 5, 12, 13, 19, 24, 31, 35, 40]. Recent approaches [1,4,8,15,28,32,36,38,39] have shifted towards integrating temporal information not only in the training loss function but also in depth prediction. The current state-of-the-art methods [8,15,38,39] adopt the cost volume generally used in stereo matching tasks to capture geometric compatibility between images. As a pioneering work, ManyDepth [38] introduces an adaptive cost volume to overcome the scale ambiguity problem in self-supervised monocular depth estimation. To enhance multi-frame feature matching in the cost volume, Depth-Former [15] incorporates attention mechanisms, replacing conventional similarity metrics with a learnable matching function. Building on these works, we also utilize a multi-frame cost volume but effectively address the misaligned feature matching problem caused by dynamic objects through probabilistic cost volume modulation.

2.2 Dynamic Objects in Static Scene Constraint

As homography warping is employed in cost volume construction and photometric reprojection loss based on the assumption of a static scene, the presence of moving objects inevitably causes incorrect matching costs and misleading supervision. To tackle the inherent challenges of multi-view inconsistency for dynamic objects in self-supervised depth learning, two key steps are essential: (1) dynamic objects should be identified from the rigid background, and (2) errors in cost volume and reprojection loss must be rectified.

Discerning dynamic objects. To identify dynamic areas, a typical approach is to use a pretrained semantic [8, 16, 19] network or an instance segmentation network [3, 4, 21, 39]. While leveraging an useful off-the-shelf network is effective in discerning moving objects, it comes with several drawbacks, including an added computational burden, the potential inclusion of static objects in segmentation masks, and confinement to predefined classes. In contrast, our proposed ProDepth identifies potential moving objects solely based on the provided images, eliminating the requirement for additional information.

Rectifying errors caused by dynamic objects. As discussed in Sec. 1, recent multi-frame based methods address the mismatch problem in the cost volume through indirect approaches, such as supervising predicted depth with single-frame depth for potential dynamic areas at the loss-level [15, 38] or removing the motion of dynamic objects at the input-level [8]. In contrast, our ProDepth aims to directly rectify the erroneous matching costs with the proposed Probabilistic Cost Volume Module (PCVM), performing motion-aware adaptive fusion of single-frame and multi-frame cues in a probabilistic manner. Additionally, to tackle the incorrect supervision in dynamic areas, existing methods [8, 12, 15, 38, 42] use a binary mask to exclude the computed loss in those

regions. However, binary masking for the estimated moving objects may not adequately consider possible dynamic areas involving ambiguous probability. Instead, we propose a loss reweighting strategy, which partially reduces incorrect supervision based on the inferred probability.

3 Method

3.1 Self-Supervised Monocular Depth Learning

Given the target image I_t and temporally adjacent source images $\{I_s \mid s \in \{t-1, t+1\}\}$, we can warp I_s to the view point of I_t with the estimated depth of the target image D_t and the relative camera pose $T_{t\to s}$:

$$I_{s \to t}(D_t) = I_s \langle \operatorname{proj}(D_t, T_{t \to s}, K) \rangle, \tag{1}$$

where K is the known camera intrinsics, $\operatorname{proj}(\cdot)$ indicates the projection of 3D points from D_t into the camera of I_s , and $\langle \cdot \rangle$ is the pixel sampling operator. For self-supervised learning of depth and camera ego-motion, the photometric reprojection loss is generally used for optimization that consists of structure similarity (SSIM) [37] and L_1 loss terms:

$$\mathcal{L}_p(D_t) = \alpha \ \frac{1 - \text{SSIM}(I_t, I_{s \to t}(D_t))}{2} + (1 - \alpha) \ \|I_t - I_{s \to t}(D_t)\|_1, \qquad (2)$$

where α is commonly set to 0.85. Importantly, this reprojection loss provides misleading supervision for dynamic areas because the image warping process is based on the static scene assumption.

3.2 Overview

The proposed architecture contains three major components that address the inconsistency issue caused by moving objects. Initially, we identify uncertainty by analyzing depth maps estimated from auxiliary depth decoders (Sec. 3.3). Subsequently, a PCVM module rectifies erroneous matching costs in the cost volume for dynamic areas by uncertainty-aware adaptive fusion of probability distributions of depth candidates from single- and multi-frame cues (Sec. 3.4). At the training-level, we mitigate misleading self-supervision devising a loss reweighting strategy (Sec. 3.5). The overall framework is summarized in Fig. 2.

3.3 Auxiliary Depth Estimations and Uncertainty Reasoning

To reason uncertainty associated with the static scene assumption and compensate for errors in dynamic areas in multi-frame cost volume, our framework incorporates two auxiliary depth estimations: single-frame depth and cost volume depth.

Probabilistic single-frame depth estimation. We employ a lightweight network, denoted as θ_{single} , to estimate single-frame depth $D_{\text{single}} \in \mathbb{R}^{H \times W}$ using



Fig. 2: Overview of the proposed ProDepth. We construct the multi-frame cost volume with I_s and I_t , and estimate single-frame depth as a Gaussian distribution using the target image I_t . In an auxiliary branch, uncertainty is inferred by comparing D_{single} and D_{cv} , where the latter is estimated from cost volume features. To rectify erroneous cost volume, a PCVM module adaptively fuses probabilities derived from single- and multi-frame cues. Furthermore, we incorporate a loss reweighting strategy in $\mathcal{L}_{up,s}$ and $\mathcal{L}_{up,s}^{log}$ to mitigate errors caused by moving objects at the training-level. Note that the probability distribution of a dynamic pixel is illustrated as an example.

a target image $I_t \in \mathbb{R}^{H \times W}$. To estimate the depth as a probability distribution, we adopt the predictive approach [18,20,30], configuring the network to output the mean μ and variance σ^2 of the distribution in the final layer. Specifically, we model the predictive distribution as a heteroscedastic Gaussian, minimizing the negative log-likelihood criterion. For supervised learning with ground-truth depth D^* , the negative log-likelihood is given by:

$$-\log p(D^*|\mu, \sigma) = \frac{(D^* - \mu)^2}{\sigma^2} + \log \sigma^2.$$
(3)

In our self-supervised learning scenario, where ground-truth D^* is unavailable, we predict the variance map $\sigma_p^2 \in \mathbb{R}^{H \times W}$ for pixel-wise photometric matching between the target image and warped image as shown in [20, 30]:

$$\mathcal{L}_p^{\log}(D_{\text{single}}) = \frac{(\mathcal{L}_p(D_{\text{single}}))^2}{\sigma_p^2} + \log \sigma_p^2.$$
(4)

Through log-likelihood maximization of \mathcal{L}_p^{\log} , we estimate single-frame depth as a probability distribution with mean D_{single} and variance σ_p^2 .



Fig. 3: The identification of dynamic objects. In contrast to the binary consistency mask generated in ManyDepth [38], our uncertainty reasons the probability of moving objects with structural awareness.

Cost volume depth estimation and uncertainty reasoning. In multiframe depth encoder ϕ_{enc} , we first encode I_t and I_s into *C*-dimensional downsampled features F_t and F_s with size $H/4 \times W/4 \times C$. We then construct the cost volume C to measure the multi-frame matching costs for hypothesized depth candidates $d = \{d_i \mid i \in \{1, 2, ..., k\}\}$. Depth candidates are perpendicular to the optic axis of I_t and uniformly sampled in log space by spatial-increasing discretization [9]:

$$d_i = e^{\log(d_1) + \frac{i}{k-1}\log(d_k/d_1)},\tag{5}$$

where depth candidates range from d_1 to d_k , representing the minimum and maximum depth values, respectively. For each depth candidate d_i , the source feature F_s is warped to the view point of I_t , producing $F_{s\to t}(d_i)$ similar to $I_{s\to t}(D_t)$ in Eq. 1. We compute the per-pixel matching costs for all d_i by the absolute L_1 difference between F_t and $F_{s\to t}$, and aggregate feature channels by average pooling to obtain the cost volume $C \in \mathbb{R}^{H/4 \times W/4 \times k}$. The cost is expected to be lower for the depth candidate that is closer to the ground-truth depth.

As the cost volume construction involves the static scene assumption in the warping process of $F_{s \to t}(d_i)$, the inconsistent geometric locations of moving objects result in misaligned feature matching with incorrectly computed matching costs. The corrupted cost distribution in dynamic areas then leads to erroneous depth estimation, degrading overall performance. However, this corruption can be leveraged to identify moving objects by comparing with accurately predicted single-frame depth. ManyDepth [38] generates a binary mask called *consistency mask* where single-frame depth D_{single} and *argmin* of the matching costs d_{low} (*i.e.*, d_i with lowest cost) differ significantly, considering it as unreliable region involving high uncertainty. The problem with this approach is that the mask relying on lowest cost cannot clearly mask out moving objects because d_{low} is computed for each pixel independently based on feature distance without an understanding of spatial correlation between pixels, *i.e.* structural awareness.

To overcome these limitations, we devise an auxiliary decoder ψ_{dec} that estimates the depth using corrupted matching costs of the cost volume. Our main

observation is that depth estimation goes beyond capturing pixel-level geometric information; it integrates structural awareness, ensuring that pixels within the same object demonstrate consistent depth values. Decoding pixel-level inconsistencies embedded in the cost volume into depth $D_{\rm cv}$ results in consistent errors within moving objects, as shown in Fig. 3. Based on the generated clear corruption, we discern the uncertainty $U \in [0, 1]^{H \times W}$ by computing the absolute difference between $D_{\rm single}$ and $D_{\rm cv}$ and normalizing it into [0, 1] with the mapping function $\mathcal{M}(a, b) = 1 - e^{-\beta |a-b|}$:

$$U = \mathcal{M}(D_{\text{single}}, D_{\text{cv}}) = 1 - e^{-\beta |D_{\text{single}} - D_{\text{cv}}|}, \tag{6}$$

where β is empirically set to 0.6. Unlike the obscure binary mask [38] generated from roughly computed per-pixel lowest cost, our uncertainty precisely indicates the probability of object-level corruption, which can be utilized as an useful cue for identifying dynamic objects.

3.4 Probabilistic Cost Volume Modulation

Contrary to existing works, we address the errors embedded in the cost volume by directly modulating the matching cost distribution. We first transform the single- and multi-frame cues into a representation of probability distribution along depth candidates and derive the modulated cost distribution for each pixel by adaptively fusing those distributions based on the uncertainty.

Single-frame depth as probability distribution. As shown in Eq. 4, we estimate the mean D_{single} and variance σ_p^2 of the single-frame depth as Gaussian distribution. By using the probability density function of Gaussian distribution $\mathcal{N}(D_{\text{single}}, \sigma_p^2)$, we can compute the probability of each depth candidate d_i given pixel x:

$$p_{\text{single}}(d_i|x) = \frac{1}{\sqrt{2\pi\sigma_p^2(x)}} \exp(-\frac{(d_i - D_{\text{single}}(x))^2}{2\sigma_p^2(x)}),\tag{7}$$

where $D_{\text{single}}(x)$ and $\sigma_p^2(x)$ indicate the estimated mean and variance values for pixel x, respectively.

Multi-frame matching costs as probability distribution. For the initially constructed cost volume $C \in \mathbb{R}^{H/4 \times W/4 \times k}$, we denote the matching cost of depth candidate d_i for the given pixel x as C(x, i). The per-pixel costs are converted into probabilities $p_{cv}(d|x)$ using the softmax function:

$$p_{\rm cv}(d_i|x) = \frac{\exp(-\mathcal{C}(x,i))}{\sum_{j=1}^k \exp(-\mathcal{C}(x,j))},\tag{8}$$

where negative costs are used for softmax because the depth candidate with lower cost holds a higher probability.

Cost volume modulation. The modulation of cost volume involves the fusion of the probability distributions of depth candidates derived from single-frame depth and multi-frame cost volume considering the uncertainty. To preserve the relative importance of each distributions after fusion, we adopt weighted

geometric mean (Eq. 9) which allows for the retention of depth candidates with highest probability due to its multiplicative nature. In contrast, the commonly used weighted arithmetic mean (weighted sum) with additive nature may not guarantee the preservation of depth candidates at the maximum due to the linear combination of distributions. We present an ablation study of the fusion strategy in the supplementary material.

Based on the computation of weighted geometric mean, probabilities $p_j \in \{p_{\text{single}}, p_{\text{cv}}\}$ are multiplied with consideration for each weight $w_j \in \{U, 1 - U\}$ to derive the fused probability distribution P(d|x) (Eq. 10):

$$P(d|x) = (\prod_{j} p_j(d|x)^{w_j})^{1/\sum_j w_j}$$
(9)

$$= p_{\text{single}}(d|x)^{U(x)} \cdot p_{\text{cv}}(d|x)^{1-U(x)}.$$
 (10)

For non-static pixels with high uncertainty, the distribution of single-frame depth p_{single} exerts a greater influence, whereas for static pixels with low uncertainty, the distribution of the cost volume p_{cv} carries more weight. The fused probability distribution P is then re-scaled to the scale of original cost volume C by min-max normalization to obtain modulated cost volume C_m :

$$\mathcal{C}_m(x,i) = \frac{\max(P(d|x)) - P(d_i|x)}{\max(P(d|x)) - \min(P(d|x))} (\max(\mathcal{C}(x)) - \min(\mathcal{C}(x))) + \min(\mathcal{C}(x)),$$
(11)

where $\frac{\max(P(d|x)) - P(d_i|x)}{\max(P(d|x)) - \min(P(d|x))}$ inverts the fused probability distribution of depth candidates while normalizing to [0, 1], since lower matching cost indicates higher probability in the original cost volume. The final multi-frame depth D_{multi} is subsequently estimated from C_m using the decoder ϕ_{dec} , wherein errors caused by dynamic objects are rectified through the cost volume modulation.

3.5 Learning without Dynamic Objects

Uncertainty-aware loss reweighting strategy. As addressed in Sec. 3.1, an optimization of photometric reprojection loss \mathcal{L}_p provides incorrect supervision for non-static pixels. To exclude the misleading loss for dynamic areas, we devise a loss reweighting strategy that adjusts the computed reprojection loss based on the uncertainty. The uncertainty-aware photometric reprojection loss \mathcal{L}_{up} is formulated as:

$$\mathcal{L}_{up} = M \odot (1 - U) \odot \mathcal{L}_p, \quad M = [U < \gamma], \tag{12}$$

where \odot is element-wise product and $[\cdot]$ denotes the Iverson bracket. The computed \mathcal{L}_p is reweighted based on the per-pixel probability of uncertainty $U \in [0,1]^{H \times W}$, and additional binary mask M is applied to rigorously exclude pixels involving high uncertainty. In comparison to conventional binary masking methods employed in existing works, our loss reweighting strategy is more effective in preventing erroneous depth overfitting for moving objects. This effectiveness

stems from its ability to partially reduce incorrect supervision for areas with ambiguous uncertainty, which may not be adequately addressed by a binary mask. Combining both binary masking and probabilistic reweighting allows us to mitigate the risk in defining learning objectives for potential dynamic areas while unequivocally excluding incorrect supervision associated with high uncertainty. **Objective functions.** Continuing with the established procedure [8, 15, 38], we incorporate the edge-aware smoothness loss \mathcal{L}_s [11] to regularize the smoothness of the predicted depth map, and consistency loss \mathcal{L}_c [8, 15, 38] to ensure multi-frame depth to be similar to single-frame depth in dynamic areas. We denote $\mathcal{L}_{up,s} = \mathcal{L}_{up} + \lambda_s \mathcal{L}_s$ as a combination of uncertainty-aware photometric reprojection loss and smoothness loss, and our final loss \mathcal{L} is

$$\mathcal{L} = \sum_{x} \left[\mathcal{L}_{up,s}(D_{\text{multi}}) + \lambda_1 \mathcal{L}_{up,s}^{log}(D_{\text{single}}) + \lambda_2 \mathcal{L}_p(D_{\text{cv}}) + \lambda_3 \mathcal{L}_c \right],$$
(13)

where x indicates pixel index. For multi- and single-frame depth estimation, our uncertainty-aware reprojection loss \mathcal{L}_{up} is employed to prevent erroneous overfitting for moving objects. In contrast, \mathcal{L}_p is used in cost volume depth estimation to encourage corruption, enabling the identification of dynamic regions with moving objects through significant depth difference between D_{cv} and D_{single} . By allowing incorrect self-supervision in dynamic areas, cost volume decoder ψ_{dec} learns to produce erroneous depth in moving objects. Note that for $\mathcal{L}_p(D_{cv})$, backpropagation is enabled exclusively for the parameters of the cost volume decoder, while the gradients are halted from flowing through the cost volume.

4 Experiments

We evaluate the performance of our approach on two challenging datasets, Cityscapes [6] and KITTI [10], recognized benchmarks for depth estimation. Since Cityscapes dataset contains more moving objects compared to KITTI, our experiments are mainly focused on Cityscapes to verify the performance improvement in dynamic scenes. We conduct quantitative and qualitative comparisons with state-of-the-art methods, and an extensive ablation study to substantiate the contributions of the proposed components. Given the importance of evaluating performance in dynamic regions for our work, additional experimental results can be found in the supplementary material.

4.1 Experimental Setup

Dataset. In our study of the Cityscapes dataset, we use a set of pre-processed 58,335 training images provided by [8], along with 1,525 images for testing. For the KITTI dataset, we adhere to the Eigen split [7] following established practices [1,8,15,38]. This split encompasses 39,810 training images, 4,424 validation images, and 697 test images. In both datasets, we exclusively use unlabeled video frames, without incorporating additional segmentation masks or optical flow information. The ground-truth depth information is employed solely for evaluation, and we constrain the predicted depth values to be below 80 meters.

Table 1: Depth evaluation on the Cityscapes and KITTI datasets. Semantics indicates

 the use of additional semantic information.

	Method	Test frames Sen	Somantics	$W \times H$	Error metric (\downarrow)			Accu	Accuracy metric (\uparrow)		
	weenod		Semantics		Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^{2}$	$\delta < 1.25^{3}$
	Struct2depth (M) [3]	1	1	416×128	0.141	1.026	5.291	0.215	0.816	0.945	0.979
	Videos in the wild [13]	1	1	416×128	0.128	0.959	5.230	0.212	0.845	0.947	0.976
	Johnston et al. [17]	1		640×192	0.111	0.941	4.817	0.189	0.885	0.961	0.981
	Packnet-SFM [14]	1		640×192	0.111	0.785	4.601	0.189	0.878	0.960	0.982
	Monodepth2 [12]	1		640×192	0.110	0.831	4.642	0.187	0.883	0.962	0.982
	HR-Depth [26]	1		640×192	0.109	0.792	4.632	0.185	0.884	0.962	0.983
E	Guizilini et al. [16]	1	1	640×192	0.102	0.698	4.381	0.178	0.896	0.964	0.984
È	Lite-Mono [42]	1		640×192	0.101	0.729	4.454	0.178	0.897	0.965	0.983
K	Patil et al. [28]	N		640×192	0.111	0.821	4.650	0.187	0.883	0.961	0.982
	Wang et al. [36]	2(-1, 0)		640×192	0.106	0.799	4.662	0.187	0.889	0.961	0.982
	ManyDepth [38]	2(-1, 0)		640×192	0.098	0.770	4.459	0.176	0.900	0.965	0.983
	DynamicDepth [8]	2(-1, 0)	1	640×192	0.096	0.720	4.458	0.175	0.897	0.964	0.984
	DepthFormer [15]	2(-1, 0)		640×192	0.090	0.661	4.149	0.175	0.905	<u>0.967</u>	0.984
	DualRefine [1]	2(-1, 0)		640×192	<u>0.090</u>	0.658	4.237	<u>0.171</u>	0.912	<u>0.967</u>	0.984
	ProDepth	2(-1, 0)		640 imes 192	0.086	0.629	4.139	0.166	0.918	0.969	0.984
	Pilzer et al. [29]	1		512×256	0.240	4.264	8.049	0.334	0.710	0.871	0.937
	Struct2Depth 2 [4]	1		416×128	0.145	1.737	7.280	0.205	0.813	0.942	0.976
	Monodepth2 [12]	1		416×128	0.129	1.569	6.876	0.187	0.849	0.957	0.983
Cityscapes	Videos in the Wild [13]	1		416×128	0.127	1.330	6.960	0.195	0.830	0.947	0.981
	Li et al. [23]	1		416×128	0.119	1.290	6.980	0.190	0.846	0.952	0.982
	Lee et al. [22]	1		832×256	0.116	1.213	6.695	0.186	0.852	0.951	0.982
	InstaDM [21]	1	1	832×256	0.111	1.158	6.437	0.182	0.868	0.961	0.983
	Struct2Depth 2 [4]	3(-1, 0, +1)	1	416×128	0.151	2.492	7.024	0.202	0.826	0.937	0.972
	ManyDepth [38]	2(-1, 0)		416×128	0.114	1.193	6.223	0.170	0.875	0.967	0.989
	DynamicDepth [8]	2(-1, 0)	1	416×128	0.103	1.000	5.867	<u>0.157</u>	0.895	0.974	0.991
	ProDepth	2(-1, 0)		416×128	0.095	0.876	5.531	0.146	0.908	0.978	0.993

Metrics. We evaluate the depth performance using widely adopted metrics [7], including four error metrics (Abs Rel, Sq Rel, RMSE, and RMSE log) and three accuracy metrics ($\delta < 1.25$, $\delta < 1.25^2$, and $\delta < 1.25^3$).

4.2 Results on Cityscapes

Table 1 presents a performance comparison between our approach and stateof-the-art methods on the Cityscapes [6] and KITTI [10] datasets. Notably, for the Cityscapes dataset, which includes a significant number of moving objects, our proposed ProDepth achieves a remarkable improvement over existing methods across all metrics. It is worth highlighting that ProDepth, relying solely on the given input images, outperforms approaches [4, 8, 21] that utilize additional semantic information. This underscores the effectiveness of our uncertainty reasoning in discerning dynamic objects. Additionally, we present qualitative results on the Cityscapes test set in Fig. 4. While related works [8, 38] exhibit relatively high estimation errors in dynamic areas, our ProDepth demonstrates superior performance.

4.3 Results on KITTI

We further evaluate our proposed ProDepth on the KITTI dataset using the Eigen split. According to the statistics analyzed in [8], the pixels indicating movable objects in dynamic classes constitute 0.34% of all pixels. As static objects are considered together in the statistics, KITTI involves a fewer dynamic areas



Fig. 4: Qualitative results on Cityscapes. Red and yellow boxes indicate moving and static objects. Error maps depict large depth errors in red and small in blue.

compared to Cityscapes. Nevertheless, our model still outperforms recent works including both single-frame and multi-frame based approaches. This demonstrates that our probabilistic fusion of single-frame and multi-frame cues also benefits the prediction in static scenes.

4.4 Ablation Study

In Table 2, we conduct extensive ablation study to evaluate three major contributions: (1) uncertainty reasoning with an auxiliary depth decoder, (2) a probabilistic cost volume modulation (PCVM) module, and (3) an uncertainty-aware loss reweighting strategy.

Uncertainty reasoning. The identification of dynamic areas can be represented in a binary or weighted (probabilistic) manner. ManyDepth [38] adopts a binary consistency mask estimated at the coarse feature-level, while DynamicDepth [8] employs a pretrained semantic segmentation network to identify the movable objects, as shown in Fig. 4. To assess the effectiveness of each mask, we substitute our uncertainty reasoning with their masks in our model during both training and inference (row $\#1\sim4$). Additionally, we convert our weighted uncertainty into a binary representation by setting a threshold (row #5,6). Our model using binary uncertainty shows similar performance to the model using segmentation masks, demonstrating that our auxiliary decoder discerns the moving objects effectively. It is noteworthy that our uncertainty does not indicate a high probability in static objects, unlike segmentation masks of DynamicDepth [8]. **PCVM.** Our proposed PCVM module performs uncertainty-aware adaptive fusion of single-frame and multi-frame cues to modulate the misaligned match-

ing cost distribution in the cost volume. The depth prediction performance is



Fig. 5: ProDepth with and without the PCVM module. Depth probability distributions of a dynamic yellow pixel are presented. Our PCVM modulates the incorrect distribution in cost volume, rectifying the errors in dynamic areas.

#	Uncertainty Reasoning		PCVM	Uncertaint	Aba Dal	Sa Dol	DMSE	DMSE log	
	Binary	Weighted	I U V M	Masking	Reweighting	Abs her	Sq Rei	RINDE	TUNDE log
1	consistency mask [38]			1		0.107	1.058	5.934	0.159
2	consistency mask [38]		1	1		0.103	0.953	5.832	0.159
3	segmentation mask [8]			1		0.100	0.961	5.620	0.150
4	segmentation mask [8]		1	1		0.101	0.965	5.647	0.150
5	$[U > \gamma]$			1		0.101	0.967	5.687	0.151
6	$[U > \gamma]$		1	1		0.099	0.944	5.616	0.151
7		U		1	1	0.100	0.964	5.630	0.151
8		U	1	1		0.098	0.903	5.551	0.148
9		U	1		1	0.097	0.894	5.512	0.146
10		U	1	1	1	0.095	0.882	5.490	0.146

 Table 2: Ablation study on the Cityscapes dataset.

enhanced by PCVM in conjunction with the consistency mask [38] and our uncertainty (row #2.6). However, when incorporating a segmentation mask, performance degrades upon adding PCVM (row #4). This decline is attributed to static objects included in segmentation masks, where only single-frame cues are utilized for those areas, while useful multi-frame cues are abandoned. Notably, with our weighted uncertainty representation, PCVM achieves a substantial performance improvement, reducing the absolute relative error from 0.100 to 0.095 (row #7,10). This demonstrates that probabilistic fusion of single-frame and multi-frame cues is more effective than selecting one of them based on the binary criterion (row #5,6). Additionally, we present qualitative results with and without the PCVM module in Fig. 5 (row #7,10). The depth candidate d_{low} with the lowest cost (*i.e.*, highest probability) for each pixel in the cost volume, and the final prediction map are provided. In the case of dynamic objects, our PCVM effectively modulates the depth probability distribution in the cost volume by integrating single-frame cues, resulting in accurate depth predictions. In contrast, the model without PCVM propagates the erroneous distribution from the cost volume to the final prediction, leading to severe errors in dynamic areas.

Uncertainty-aware loss reweighting strategy. To address the incorrect self-supervision in dynamic areas, existing methods use a binary mask to exclude computed losses in those areas. In contrast, we propose a loss reweighting strategy, which reduces the computed loss based on the inferred probability (row #9). Comparing our reweighting approach (row #9) with binary masking (row #8), we observe improved performance with the reweighting strategy. This is because binary masking lacks consideration of the detailed probability of areas being dynamic, and areas with ambiguous probability may not be adequately addressed due to thresholding. Furthermore, combining both masking and reweighting strategies yields the best performance (row #10). This approach effectively excludes incorrect supervision in areas with high uncertainty through binary masking, while partially reducing the risk of providing incorrect supervision in remaining but potential dynamic areas according to the probability.

4.5 Generalization Study

We further validate the generalization ability of the proposed ProDepth and related works [8,38] on the Waymo Open dataset [34], which encompasses numerous dynamic objects and challenging scenes like low-light conditions during nighttime. We use 202 test video sequences for evaluation. The models are pretrained on the Cityscapes dataset. Given that DynamicDepth [8] necessitates a pretrained semantic segmentation network during inference, we pre-compute masks using EfficientPS [27], which is utilized in experiments on Cityscapes and KITTI datasets. As shown in Table 3, ProDepth achieves superior performance compared to related works, showcasing its effective generalization ability. For evaluation only on dynamic areas, please refer to the supplementary material.

Table 3: Generalization study on the Waymo Open dataset.

Method	Test frames Seman	ntics Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^{3}$
ManyDepth [38]	2(-1, 0)	0.260	3.916	10.463	0.313	0.606	0.856	0.941
DynamicDepth [8]	2 (-1, 0) 🗸	0.255	3.521	9.902	0.313	0.601	0.856	0.942
ProDepth	2(-1, 0)	0.247	3.462	9.544	0.300	0.628	0.873	0.949

5 Conclusion

We present ProDepth, a multi-frame depth estimation framework addressing the inconsistency problem caused by dynamic objects in a probabilistic manner. Our contributions involve novel approaches of discerning the probability of areas being dynamic, direct rectification of misaligned cost volume with adaptive fusion of single-frame and multi-frame cues, and alleviating incorrect self-supervision in potential dynamic areas with a loss reweight strategy. ProDepth achieves state-of-the-art performance on both Cityscapes and KITTI datasets, and extensive experiments demonstrate the effectiveness of the proposed method.

Acknowledgement. This work was supported by the Yonsei Signature Research Cluster Program of 2024 (2024-22-0161) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2021-0-02068, Artificial Intelligence Innovation Hub).

References

- Bangunharcana, A., Magd, A., Kim, K.S.: Dualrefine: Self-supervised depth and pose estimation through iterative epipolar sampling and refinement toward equilibrium. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 726–738 (2023)
- Bian, J., Li, Z., Wang, N., Zhan, H., Shen, C., Cheng, M.M., Reid, I.: Unsupervised scale-consistent depth and ego-motion learning from monocular video. Advances in neural information processing systems **32** (2019)
- Casser, V., Pirk, S., Mahjourian, R., Angelova, A.: Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 8001–8008 (2019)
- Casser, V., Pirk, S., Mahjourian, R., Angelova, A.: Unsupervised monocular depth and ego-motion learning with structure and semantics. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
- Chen, Y., Schmid, C., Sminchisescu, C.: Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7063– 7072 (2019)
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: Proceedings of the IEEE international conference on computer vision. pp. 2650–2658 (2015)
- Feng, Z., Yang, L., Jing, L., Wang, H., Tian, Y., Li, B.: Disentangling object motion and occlusion for unsupervised multi-frame monocular depth. In: European Conference on Computer Vision. pp. 228–244. Springer (2022)
- Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2002–2011 (2018)
- Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 3354–3361. IEEE (2012)
- Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 270–279 (2017)
- Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into selfsupervised monocular depth estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3828–3838 (2019)

- 16 S. Woo et al.
- Gordon, A., Li, H., Jonschkowski, R., Angelova, A.: Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8977–8986 (2019)
- 14. Guizilini, V., Ambrus, R., Pillai, S., Raventos, A., Gaidon, A.: 3D packing for self-supervised monocular depth estimation. In: CVPR (2020)
- Guizilini, V., Ambruş, R., Chen, D., Zakharov, S., Gaidon, A.: Multi-frame selfsupervised depth with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 160–170 (2022)
- Guizilini, V., Hou, R., Li, J., Ambrus, R., Gaidon, A.: Semantically-guided representation learning for self-supervised monocular depth. In: ICLR (2020)
- 17. Johnston, A., Carneiro, G.: Self-supervised monocular trained depth estimation using self-attention and discrete disparity volume. In: CVPR (2020)
- Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? Advances in neural information processing systems **30** (2017)
- Klingner, M., Termöhlen, J.A., Mikolajczyk, J., Fingscheidt, T.: Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In: European Conference on Computer Vision. pp. 582–600. Springer (2020)
- Klodt, M., Vedaldi, A.: Supervising the new with the old: learning sfm from sfm. In: Proceedings of the European conference on computer vision (ECCV). pp. 698–713 (2018)
- Lee, S., Im, S., Lin, S., Kweon, I.S.: Learning monocular depth in dynamic scenes via instance-aware projection consistency. In: 35th AAAI Conference on Artificial Intelligence/33rd Conference on Innovative Applications of Artificial Intelligence/11th Symposium on Educational Advances in Artificial Intelligence. pp. 1863–1872. ASSOC ADVANCEMENT ARTIFICIAL INTELLIGENCE (2021)
- Lee, S., Rameau, F., Pan, F., Kweon, I.S.: Attentive and contrastive learning for joint depth and motion field estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4862–4871 (2021)
- Li, H., Gordon, A., Zhao, H., Casser, V., Angelova, A.: Unsupervised monocular depth learning in dynamic scenes. In: CoRL (2020)
- Li, H., Gordon, A., Zhao, H., Casser, V., Angelova, A.: Unsupervised monocular depth learning in dynamic scenes. In: Conference on Robot Learning. pp. 1908– 1917. PMLR (2021)
- Li, R., Gong, D., Yin, W., Chen, H., Zhu, Y., Wang, K., Chen, X., Sun, J., Zhang, Y.: Learning to fuse monocular and multi-view cues for multi-frame depth estimation in dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21539–21548 (2023)
- Lyu, X., Liu, L., Wang, M., Kong, X., Liu, L., Liu, Y., Chen, X., Yuan, Y.: Hrdepth: High resolution self-supervised monocular depth estimation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 2294–2301 (2021)
- Mohan, R., Valada, A.: Efficientps: Efficient panoptic segmentation. International Journal of Computer Vision 129(5), 1551–1579 (2021)
- Patil, V., Van Gansbeke, W., Dai, D., Van Gool, L.: Don't forget the past: Recurrent depth estimation from monocular video. IEEE Robotics and Automation Letters 5(4), 6813–6820 (2020)
- Pilzer, A., Xu, D., Puscas, M.M., Ricci, E., Sebe, N.: Unsupervised adversarial depth estimation using cycled generative networks. In: 3DV (2018)

- Poggi, M., Aleotti, F., Tosi, F., Mattoccia, S.: On the uncertainty of self-supervised monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3227–3237 (2020)
- Ranjan, A., Jampani, V., Kim, K., Sun, D., Wulff, J., Black, M.J.: Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In: CVPR (2019)
- Ruhkamp, P., Gao, D., Chen, H., Navab, N., Busam, B.: Attention meets geometry: Geometry guided spatial-temporal attention for consistent self-supervised monocular depth estimation. In: 2021 International Conference on 3D Vision (3DV). pp. 837–847. IEEE (2021)
- Shu, C., Yu, K., Duan, Z., Yang, K.: Feature-metric loss for self-supervised learning of depth and egomotion. In: European Conference on Computer Vision. pp. 572– 588. Springer (2020)
- Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2446–2454 (2020)
- Vijayanarasimhan, S., Ricco, S., Schmid, C., Sukthankar, R., Fragkiadaki, K.: Sfmnet: Learning of structure and motion from video. arXiv preprint arXiv:1704.07804 (2017)
- 36. Wang, J., Zhang, G., Wu, Z., Li, X., Liu, L.: Self-supervised joint learning framework of depth estimation via implicit cues. arXiv:2006.09876 (2020)
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing 13(4), 600–612 (2004)
- Watson, J., Mac Aodha, O., Prisacariu, V., Brostow, G., Firman, M.: The temporal opportunist: Self-supervised multi-frame monocular depth. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1164– 1174 (2021)
- Wimbauer, F., Yang, N., Von Stumberg, L., Zeller, N., Cremers, D.: Monorec: Semisupervised dense reconstruction in dynamic environments from a single moving camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6112–6122 (2021)
- 40. Yin, Z., Shi, J.: GeoNet: Unsupervised learning of dense depth, optical flow and camera pose. In: CVPR (2018)
- 41. Zhan, H., Garg, R., Weerasekera, C.S., Li, K., Agarwal, H., Reid, I.: Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In: CVPR (2018)
- Zhang, N., Nex, F., Vosselman, G., Kerle, N.: Lite-mono: A lightweight cnn and transformer architecture for self-supervised monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18537–18546 (2023)
- 43. Zhou, T., Brown, M., Snavely, N., Lowe, D.: Unsupervised learning of depth and ego-motion from video. In: CVPR (2017)