

From-Ground-To-Objects: Coarse-to-Fine Self-supervised Monocular Depth Estimation of Dynamic Objects with Ground Contact Prior

Jaeho Moon Juan Luis Gonzalez Bello Byeongjun Kwon Munchurl Kim

Korea Advanced Institute of Science and Technology

{jaeho.moon, juanluisgb, kbj2738, mkimee}@kaist.ac.kr

https://kaist-viclab.github.io/From_Ground_To_Objects_site/

Abstract

Self-supervised monocular depth estimation (DE) is an approach to learning depth without costly depth ground truths. However, it often struggles with moving objects that violate the static scene assumption during training. To address this issue, we introduce a coarse-to-fine training strategy leveraging the ground contacting prior based on the observation that most moving objects in outdoor scenes contact the ground. In the coarse training stage, we exclude the objects in dynamic classes from the reprojection loss calculation to avoid inaccurate depth learning. To provide precise supervision on the depth of the objects, we present a novel Ground-contacting-prior Disparity Smoothness Loss (GDS-Loss) that encourages a DE network to align the depth of the objects with their ground-contacting points. Subsequently, in the fine training stage, we refine the DE network to learn the detailed depth of the objects from the reprojection loss, while ensuring accurate DE on the moving object regions by employing our regularization loss with a cost-volume-based weighting factor. Our overall coarse-to-fine training strategy can easily be integrated with existing DE methods without any modifications, significantly enhancing DE performance on challenging Cityscapes and KITTI datasets, especially in the moving object regions.

1. Introduction

Recent neural-network-based methods for depth estimation (DE) from 2D images have shown promising performance as the demand for 3D information grows with navigation, robotics, and AR/VR applications. While supervised methods [2, 10, 29, 33, 45] utilize expensive and sparse ground truth (GT) depth for training, self-supervised methods learn depth by minimizing photometric errors (or reprojection loss) between the target frame and the warped frames from stereo images [11, 12, 14, 15, 48, 50] or consecutive frames in monocular videos [13, 17, 51, 54].

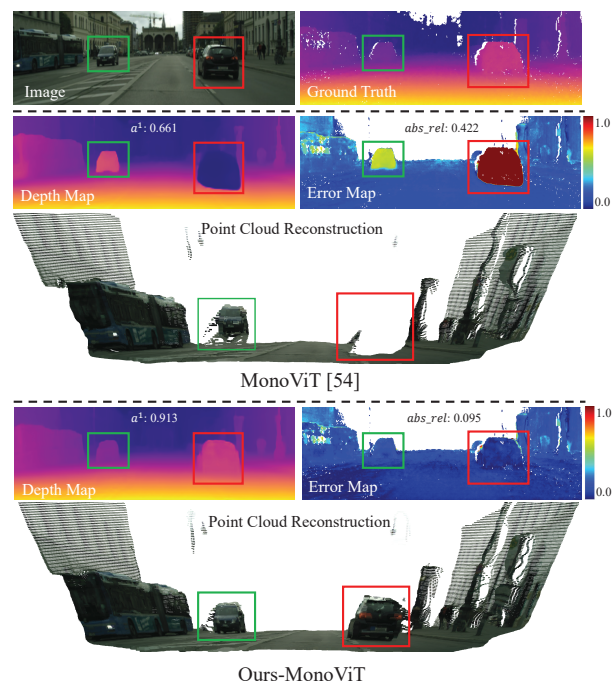


Figure 1. The effect of our coarse-to-fine self-supervised DE training strategy applied to MonoViT [54]. In the error maps, blue indicates small errors, and red indicates large errors.

In the pipeline of self-supervised depth learning for monocular videos, a DE network often learns inaccurate depth of moving objects that break the static scene assumption but still minimize the reprojection loss for self-supervision [13, 17, 18, 27]. To handle the moving object problem, the ‘automasking’ [13] is proposed to exclude the static pixels in sequential frames from the reprojection loss calculation. More recent work suggested to estimate 3D object motions [22, 30–32] or to disentangle object motion for cost volume construction [9]. However, these methods confront inherent ambiguities in learning the depth and motion of the objects without precise supervision.

In this paper, we introduce a ground contacting prior to effectively resolve the depth ambiguity of moving objects in a self-supervised manner. The ground contacting prior is based on an observation that most objects classified as dynamic in outdoor scenes, such as cars, bicycles, or pedestrians, invariably tend to make contact with the ground, thereby sharing similar depth at their ground contact points.

Leveraging this crucial insight, we introduce a coarse-to-fine training strategy for self-supervised monocular DE. In the initial coarse training stage, we train a DE network by employing the novel Ground-contacting-prior Disparity Smoothness Loss (GDS-Loss) that aligns the depth of objects in dynamic classes with the depth of their ground contact points. Also, we exclude the objects from the reprojection loss computation to avoid learning inaccuracies in DE. We further fine-tune the DE network in the fine training stage for learning to capture detailed depth over the surfaces of objects in dynamic classes by employing the reprojection loss for the whole image region. We introduce a regularization loss that allows the DE network to learn the detailed depth of non-moving objects by the reprojection loss while avoiding inaccurate learning of DE on the moving object regions with a cost-volume-based weighting factor.

As shown in Fig. 1, MonoViT [54] predicts inaccurate depth that make cars float (green box) or sunken under the ground (red box) in the reconstructed point cloud. On the other hand, MonoViT trained with our coarse-to-fine training strategy, denoted as Ours-MonoViT, predicts the accurate depth of the cars, so that they are appropriately standing on the ground in the reconstructed point cloud. In summary, our contributions are three-fold:

- We, for the first time, propose to utilize the ground contacting prior as a self-supervision for DE on objects in dynamic classes, presenting the novel Ground-contacting-prior Disparity Smoothness Loss (GDS-Loss).
- We introduce a regularization loss with a cost-volume-based weighting factor, which allows fine-tuning with the reprojection loss while ensuring adherence to the ground contacting prior on the moving object regions.
- Our coarse-to-fine training strategy can easily be integrated into the existing DE networks, boosting their DE performance to achieve state-of-the-art results on the challenging Cityscapes and KITTI datasets.

2. Related Work

2.1. Self-supervised Monocular Depth Estimation

Since measuring the GT depth is expensive, especially in outdoor environments, self-supervised learning of DE has been intensively studied [7, 11–13, 15, 20, 40, 43, 47, 51, 55]. Without the GT depth, neural networks (NN) have been adopted to learn depth from stereo images [14, 41, 42, 50,

57] or sequential video frames [4, 18, 25, 46, 53] by minimizing the photometric errors, or reprojection loss, between a target view and the warped reference view based on the predicted depth. Since the usage of stereo images inherently has limitations such as requirements of synchronized binocular cameras and calibration, self-supervised learning from monocular videos for DE has been actively investigated.

From monocular videos, NN-based methods learn from structure-from-motion techniques. Zhou *et al.* [55] incorporated two separate networks to learn camera ego-motion and depth simultaneously from monocular videos. Godard *et al.* proposed a CNN-based architecture, Monodepth2 [13], which handles occlusion problems based on the minimum reprojection loss. Various approaches have been explored: ranging from improving network architectures, [1, 17, 24, 37, 52, 54], using semantics-guidance [25, 27], to employing sequential frames in test time by utilizing the cost volumes [9, 19, 51].

Recently, Petrovai *et al.* [40] suggested a two-stage training pipeline to transfer the knowledge of the high-resolution DE as explicit pseudo labels into the low-resolution DE training of a student network. This process involves filtering out noisy pseudo depth labels by assessing 3D consistency across sequential frames. In contrast, our coarse-to-fine training strategy guides a DE network to learn the depth of objects in dynamic classes based on the ground contacting prior (GDS-Loss) in the coarse stage and to further refine them in the fine stage. Our approach employs a regularization loss in the fine stage to regularize the reprojection loss for enhancing the detailed depth predictions in static object regions while encouraging the precise learning of DE on the moving object regions with a cost-volume-based weighting factor.

2.2. Handling Moving Objects

Since the self-supervised monocular DE pipelines learn depth and camera ego-motion from consecutive frames under the rigid scene assumption, independently moving objects induce them to learn erroneous depth. To overcome this problem, Monodepth2 [13] utilized automasking that excludes pixels static in sequential frames, but it is only capable of masking out the objects moving at the same speed and in the same direction as the camera from the calculation of the reprojection loss. Although other approaches suggested filtering out the images with moving objects from training scenes [18], and masking out the detected moving objects in photometric loss computation [27], discerning whether an object is moving or not involves implicit uncertainty, which may lead to inaccurate DE performance.

Recent work [3, 9, 16, 22, 30–32] attempted to learn depth by decoupling object motion from camera ego-motion, exploiting a regularization term for 3D translation fields [32], or predicting individual motions of objects

by utilizing off-the-shelf segmentation models [3, 30, 31]. Feng *et al.* [9] suggested an occlusion-aware cost volume for multi-frame DE by adopting a cycle-consistency loss. However, without precise supervision of depth or motion of moving objects, these methods often yield inferior results, since the reprojection loss can be minimized even with inaccurate depth and object motion predictions. In contrast, our proposed GDS-Loss provides precise depth supervision on the objects classified as dynamic classes, regardless of their moving speed and directions based on the ground point they stand, which is ground contacting prior. Moreover, our method can easily be integrated with existing DE methods by utilizing instance segmentation masks only during training without necessitating of training an auxiliary object motion estimation network.

3. Method

3.1. Self-supervised Monocular Depth Estimation

In this subsection, we briefly describe the conventional self-supervised monocular depth estimation (DE) pipeline in a self-contained manner. Given two consequent frames I_t and $I_{t'}$ at time t and t' , respectively, from a monocular video, a DE network θ_{depth} predicts a depth map D_t for target image I_t , as given by

$$D_t = \theta_{depth}(I_t). \quad (1)$$

Also, a pose estimation network θ_{pose} predicts a camera ego-motion $T_{t \rightarrow t'}$ from I_t to $I_{t'}$. We can obtain a synthesized frame $I_{t'}^w$ by warping $I_{t'}$ with the projected coordinates based on the estimated depth D_t , pose $T_{t \rightarrow t'}$ and the camera intrinsic matrix K , as given by

$$I_{t'}^w = DW(I_{t'}, proj(D_t, T_{t \rightarrow t'}, K)), \quad (2)$$

where $proj$ denotes the projection of the 3D transformed points and DW indicates the differentiable warping process [23]. Then the reprojection loss L_{rep} defined as the photometric errors between $I_{t'}^w$ and I_t , consisting of L1 and SSIM [49] error terms weighted by α , as given by

$$L_{rep} = \alpha \cdot L1(I_t, I_{t'}^w) + (1 - \alpha) \cdot SSIM(I_t, I_{t'}^w). \quad (3)$$

In addition, an edge-aware smoothness loss L_{sm} [12] on the disparity d_t (inverse of depth D_t) aims to encourage locally smooth depth maps while preserving the edges in the images, as is given by

$$L_{sm} = |\partial_x \hat{d}_t| e^{-|\partial_x I_t|} + |\partial_y \hat{d}_t| e^{-|\partial_y I_t|}, \quad (4)$$

where \hat{d}_t is the normalized disparity by the mean of d_t . Specifically, each term in Eq. 4 consists of the disparity gradient exponentially weighted by the negative horizontal or vertical gradient of I_t . As the color gradient $\partial_x I_t$ (or $\partial_y I_t$) decreases, L_{sm} gets increased which guides θ_{depth} to

predict smooth depth in homogeneous regions. On the other hand, as $\partial_x I_t$ (or $\partial_y I_t$) increases, L_{sm} decreases to induce abrupt boundaries in the depth map. Through the training, θ_{depth} and θ_{pose} learn DE and camera ego-motion, respectively, by minimizing both L_{rep} and L_{sm} , as given by

$$L_{total} = L_{rep} + L_{sm}. \quad (5)$$

3.2. Proposed Self-supervised DE Pipeline

Fig. 2 illustrates our coarse-to-fine training strategy for self-supervised monocular DE. As depicted in the ‘Coarse Training Stage’ (also in Sec. 3.2.1), DE network θ_{depth}^1 learns the depth of objects in dynamic classes from our proposed Ground-contacting-prior Disparity Smoothness Loss L_{GDS} (Eq. 9) that provides precise supervision. The region of objects in dynamic classes are masked out from the calculation of L_{rep} (Eq. 6) to avoid inaccurate learning of DE. After the ‘Coarse Training Stage’, the trained DE network θ_{depth}^1 and pose estimation network θ_{pose} are fixed and denoted as θ_{depth}^{1*} and θ_{pose}^* , respectively. In the ‘Fine Training Stage’ (also in Sec. 3.2.2), θ_{depth}^2 , initialized as the trained θ_{depth}^1 , is further finetuned to capture detailed depth over the surfaces of objects in dynamic classes by applying an unmasked L_{rep} on the whole image. To prevent θ_{depth}^2 from learning inaccurate DE on moving object regions, we introduce our regularization loss L_{REG} (Eq. 14) with a cost-volume-based weighting factor λ_{cv} (Eq. 13).

3.2.1 Coarse Training Stage with Ground Contacting Prior

Excluding objects in dynamic classes from L_{rep} calculation. From the conventional self-supervised monocular DE pipeline in Sec. 3.1, camera ego-motion is only considered for calculating L_{rep} . Even though objects have independent motions that often differ from the camera, L_{rep} can still be minimized via the estimated camera motion. As a result, the DE network often fails to predict the precise depth of the objects, as shown in Fig. 1. In this coarse training stage, instead of adopting erroneous supervision from L_{rep} , we mask objects that belong to dynamic classes, referred to as objects in dynamic classes, such as cars, bicycles, and pedestrians from calculating L_{rep} whether the objects are moving or not. Using a pre-computed instance segmentation map for input image I_t , we define the binary dynamic instance mask as M_t , valued 1 for the region of objects in dynamic classes and 0 otherwise. The per-pixel masked reprojection loss L_{rep}^m is defined as

$$L_{rep}^m = (1 - M_t)L_{rep}, \quad (6)$$

so that the objects in dynamic classes can effectively be excluded from the calculation of L_{rep} , as also depicted in the ‘Coarse Training Stage’ of Fig. 2.




Figure 2. Overview of our coarse-to-fine training strategy. In the coarse training stage, a DE network learns the depth of static region from the masked reprojection loss L_{rep}^m and the depth of the objects in dynamic classes from our GDS-Loss L_{GDS} . In the fine training stage, we further refine the DE network with unmasked L_{rep} while regularizing to ensure consistent depth prediction with the reliable DE network trained in the coarse stage by employing our regularization loss L_{REG} with a cost volume-based weighting factor λ_{cv} .

Leveraging the ground contacting prior for supervision. A predominant cause of imprecise depth prediction for moving objects is insufficient supervision in the self-supervised learning of DE. To address this issue, we introduce the ground contacting prior, rooted in the observation that most objects in dynamic classes such as cars, bicycles, and pedestrians, are in contact with the ground, resulting in similar depth with the depth of ground point they touch. Based on this, we introduce a novel Ground-contacting-prior Disparity Smoothness Loss (GDS-Loss) that can precisely supervise the depth of objects in dynamic classes. Our GDS-Loss is designed to guide DE networks for consistent depth predictions between the objects and their contacting ground locations. As the ground is static, the DE networks can accurately predict the depth of the ground, which can be utilized for strong self-supervision on the objects.

Ground-contacting-prior Disparity Smoothness Loss. Our GDS-Loss is implemented by extending the conventional edge-aware disparity smoothness loss (Eq. 4). Our GDS-Loss focuses on the vertical gradient of disparity $|\partial_y \hat{d}_t|$ to induce consistency of DE in the region between objects in dynamic classes and their contacting ground. Formally, for $\hat{d}_t \in \mathbb{R}^{n \times m}$, $|\partial_y \hat{d}_t|$ represents the disparity difference between two vertically adjacent pixel locations, (i, j) and its bottom neighbor $(i + 1, j)$, which is expressed as:

$$|\partial_y \hat{d}_t|(i, j) = |\hat{d}_t(i, j) - \hat{d}_t(i + 1, j)|, \quad (7)$$

where $|\partial_y \hat{d}_t| \in \mathbb{R}^{(n-1) \times m}$. In order to enforce consistency of depth between the objects in dynamic classes and their ground points of contact, we introduce a ground-contacting-prior mask M_{gr} given as

$$M_{gr}(i, j) = \gamma \cdot M_t(i, j) + (1 - M_t(i, j)) \quad (8)$$

where $M_{gr} \in \mathbb{R}^{(n-1) \times m}$ and γ is a weighting parameter for M_{gr} , empirically setting to 100. For the location (i, j) in the object region, we let $|\partial_y \hat{d}_t|$ be highly weighted by $M_{gr}(i, j) = \gamma$, so that the disparities are enforced to be consistent with their neighboring pixels in a downward direction. On the other hand, $|\partial_y \hat{d}_t|$ is weighted the same as in Eq. 4 in the background region because of $M_{gr}(i, j) = 1$. We additionally modify the color gradients $|\partial_y I_t|$ and $|\partial_x I_t|$ to guide consistent depth prediction inside the regions of objects in dynamic classes by replacing I_t with $I_t^m = (1 - M_t)I_t$ to remove the effect from the color edges. Our GDS-Loss is defined as:

$$L_{GDS} = |\partial_x \hat{d}_t| e^{-|\partial_x I_t^m|} + |\partial_y \hat{d}_t| M_{gr} e^{-|\partial_y I_t^m|}, \quad (9)$$

so that it induces smooth depth inside the object regions while aligning them with their contacting ground points. Note that only $|\partial_y \hat{d}_t|$ is weighted by M_{gr} .

Finally, our total loss function L_{total}^1 in the coarse training stage with the masked reprojection loss and GDS-Loss is now defined by

$$L_{total}^1 = L_{rep}^m + \beta L_{GDS}. \quad (10)$$




Figure 3. Our fine training stage further refines to capture the detailed depth of the objects in dynamic classes.

3.2.2 Fine Training Stage

Refinement of the depth estimation network. The initial DE network θ_{depth}^1 , trained using the GDS-Loss (Eq. 9) in the coarse training stage, effectively aligns the depth of the objects in dynamic classes with their ground contact points. Nonetheless, it may struggle to capture the detailed depth of objects when their surfaces vary vertically or if their bottom parts are occluded by other (closer to the camera) objects, leading to homogeneous depth prediction on both objects. Thus, we introduce a fine training stage for θ_{depth}^2 to learn the complex depth of the objects in dynamic classes by applying an L_{rep} (Eq. 3) for the whole image region. Note that θ_{depth}^2 is initialized as the trained θ_{depth}^1 weights and the fixed θ_{depth}^1 provides moderate supervision in this fine training stage. Despite L_{rep} is critical to induce DE networks to learn the detailed depth of non-moving objects, such as parked cars, that follow the rigid scene assumption, it is prone to guide θ_{depth}^2 to learn the inaccurate DE on moving object regions, as discussed in Sec. 3.2.1.

Regularization with respect to θ_{depth}^{1*} . To regularize θ_{depth}^2 from learning inaccurate DE, we propose to guide θ_{depth}^2 to predict consistent depth with reliable depth predictions of θ_{depth}^{1*} that is trained based on the ground contacting prior. We introduce a per-pixel depth regularization loss that guides θ_{depth}^2 to learn detailed depth from L_{rep} while maintaining alignment with θ_{depth}^{1*} , which is expressed by

$$L_{REG}^0 = \max(|D_t^1 - D_t^2|, \delta_D), \quad (11)$$

where D_t^1 and D_t^2 represent the depth predictions of θ_{depth}^{1*} and θ_{depth}^2 for target frame I_t , respectively. δ_D is a hyperparameter for an allowable depth difference between θ_{depth}^2 and θ_{depth}^{1*} . Note that L_{REG}^0 is activated only when the depth difference $|D_t^1 - D_t^2|$ exceeds δ_D , while allowing θ_{depth}^2 to learn fine details without regularization when $|D_t^1 - D_t^2|$ is minor. However, our experiments indicate that the vanilla L_{REG}^0 fails to regularize when current $|D_t^1 - D_t^2|$ is small but quickly becomes larger by L_{rep} in the moving object regions, so that θ_{depth}^2 learns inaccurate DE.

Cost-volume-based weighting factor. To enhance regularization in the moving object regions, we introduce a

cost-volume-based approach that can identify the region by analyzing the depth discrepancy between D_t^1 and a cost-volume-induced depth. This approach is feasible because of the reliable depth prediction D_t^1 induced by the ground contacting prior. The multi-view cost volume is constructed by the L1 distances between the target frame I_t and the warped neighbor frame I'_t to the position of I_t with the depth candidates ranging from D_{min} to D_{max} , following [51]. From the cost volume, the depth for each pixel that minimizes the L1 distance among depth candidates, which is the cost-volume-induced depth D_{cv} , is expressed as:

$$D_t^{cv} = \arg \min_{D_k \in \{1, \dots, 32\}} (|DW(I'_t, \text{proj}(D_k, T_{t \rightarrow t'}, K)) - I_t|), \quad (12)$$

where D_k represents a k -th quantized depth candidate in the range of $[D_1 = D_{min}, D_{32} = D_{max}]$. D_{min} and D_{max} are determined by the minimum and maximum values of D_t^1 , respectively. By leveraging the difference between the cost-volume-induced depth D_t^{cv} and D_t^1 , we define a pixel-wise cost-volume-based-weighting factor as

$$\lambda_{cv} = \max\left(\frac{|D_t^1 - D_t^{cv}|}{\delta_D}, 1\right), \quad (13)$$

where λ_{cv} is equal to the normalized depth difference by δ_D , being always greater than or equal to 1. In order to regularize when $|D_t^1 - D_t^{cv}|$ is larger than δ_D in the moving object regions regardless of their current depth predictions, we additionally mask δ_D with $\mu_{cv} = [\lambda_{cv} = 1]$ where $[\cdot]$ is an Iverson bracket. By employing λ_{cv} and μ_{cv} into L_{REG}^0 in Eq. 11, our proposed pixel-wise regularization loss is defined as

$$L_{REG} = \lambda_{cv} \cdot \max(|D_t^1 - D_t^2|, \mu_{cv} \delta_D), \quad (14)$$

where δ_D is empirically set to 5% of D_{max} in our experiments. The regions of moving objects have large $|D_t^1 - D_t^{cv}|$ values, so L_{REG} regularizes the θ_{depth}^2 learning toward maintaining the consistency between D_t^1 and D_t^2 . In contrast, the static regions which often have small values of $|D_t^1 - D_t^{cv}|$ and $|D_t^1 - D_t^2|$ can learn detailed depth without the regularization. Since λ_{cv} penalizes more on the region that differs a lot between D_{cv} and the reliable D_t^1 , L_{REG} is a more stable regularization than the approach that directly mask out pixels determined as moving from L_{rep} . Our total loss L_{total}^2 in the fine training stage is defined as,

$$L_{total}^2 = L_{rep} + \rho L_{REG}. \quad (15)$$

Note L_{total}^1 and L_{total}^2 are averaged over all image pixels and batches for training. As shown in Fig. 3, MonoViT [54] trained with the coarse training stage predicts consistent depth between the bus and the front car in the red box, so the bus has a higher error than the car. On the other hand, our fine training stage guides Our-MonoViT to predict the accurate depth of the occluded bus with low error.

D	Methods	O.N.	Sem	Resolutions	$abs\ rel \downarrow$	$sq\ rel \downarrow$	$rmse \downarrow$	$log_{rmse} \downarrow$	$a^1 \uparrow$	$a^2 \uparrow$	$a^3 \uparrow$
Cityscapes	Struct2Depth [3]	✓	Tr	128 × 416	0.145	1.737	7.280	0.205	0.813	0.942	0.976
	Gordon et al. [16]	✓	Tr	128 × 416	0.127	1.330	6.960	0.195	0.830	0.947	0.981
	Li et al. [32]	✓		128 × 416	0.119	1.290	6.980	0.190	0.846	0.952	0.982
	Monodepth2 [13]			128 × 416	0.129	1.569	6.876	0.187	0.849	0.957	0.983
	Ours-Monodepth2		Tr	128 × 416	0.110	1.179	6.390	0.169	0.881	0.968	0.989
	DynamicDepth† [9]		Tr/Te	128 × 416	0.103	1.000	5.867	0.157	0.895	0.974	0.991
	MonoViT [54]			128 × 416	0.114	1.238	6.589	0.174	0.860	0.965	0.990
	Ours-MonoViT		Tr	128 × 416	0.096	0.930	5.806	0.152	0.905	0.976	0.992
	Lee et al. [31]	✓	Tr	256 × 832	0.116	1.214	6.695	0.186	0.852	0.951	0.982
	InstaDM [30]	✓	Tr	256 × 832	0.111	1.158	6.437	0.182	0.868	0.961	0.983
	Monodepth2 [13]			192 × 640	0.125	1.474	6.688	0.180	0.865	0.964	0.988
	Ours-Monodepth2		Tr	192 × 640	0.102	1.024	6.015	0.159	0.896	0.973	0.990
	HR-Depth [37]			192 × 640	0.120	1.253	6.714	0.179	0.857	0.963	0.988
Ours-HR-Depth		Tr	192 × 640	0.100	1.010	5.998	0.157	0.896	0.974	0.991	
RM-Depth [22]	✓		192 × 640	0.100	0.839	5.774	0.154	0.895	0.976	0.993	
CADDepth [52]			192 × 640	0.124	1.278	6.771	0.183	0.862	0.962	0.986	
Ours-CADDepth		Tr	192 × 640	0.097	0.966	5.646	0.150	0.907	0.978	0.992	
MonoViT [54]			192 × 640	0.106	1.098	6.071	0.160	0.881	0.974	0.991	
Ours-MonoViT		Tr	192 × 640	0.088	0.795	5.368	0.140	0.920	0.981	0.994	
KITTI	Struct2Depth [3]	✓	Tr	192 × 640	0.141	1.026	5.291	0.215	0.816	0.945	0.979
	Li et al. [32]	✓		192 × 640	0.130	0.950	5.138	0.209	0.843	0.948	0.978
	Gordon et al. [16]	✓	Tr	192 × 640	0.128	0.959	5.230	0.212	0.845	0.947	0.976
	SGDepth [27]		Tr	192 × 640	0.117	0.907	4.844	0.196	0.875	0.958	0.980
	Lee et al. [31]	✓	Tr	256 × 832	0.114	0.876	4.715	0.191	0.872	0.955	0.981
	InstaDM [30]	✓	Tr	256 × 832	0.112	0.777	4.772	0.191	0.872	0.959	0.982
	Monodepth2 [13]			192 × 640	0.115	0.903	4.863	0.193	0.877	0.959	0.981
	Ours-Monodepth2		Tr	192 × 640	0.112	0.866	4.766	0.190	0.879	0.960	0.982
	PackNet [17]			192 × 640	0.111	0.785	4.601	0.189	0.878	0.960	0.982
	Johnston et al. [24]			192 × 640	0.111	0.941	4.817	0.189	0.885	0.961	0.981
	HR-Depth [37]			192 × 640	0.109	0.792	4.632	0.185	0.884	0.962	0.983
	Ours-HR-Depth		Tr	192 × 640	0.108	0.775	4.614	0.184	0.886	0.962	0.983
	FSRE-Depth [25]		Tr	192 × 640	0.105	0.708	4.546	0.182	0.886	0.964	0.983
	Petrovai et al. [40]			192 × 640	0.106	0.751	4.485	0.180	0.885	0.964	0.984
	RM-Depth [22]	✓		192 × 640	0.108	0.710	4.513	0.183	0.883	0.964	0.983
	CADDepth [52]			192 × 640	0.105	0.769	4.535	0.181	0.892	0.964	0.983
	Ours-CADDepth		Tr	192 × 640	0.103	0.730	4.427	0.179	0.895	0.966	0.984
DynamicDepth† [9]		Tr/Te	192 × 640	0.096	0.720	4.458	0.175	0.897	0.964	0.984	
MonoViT [54]			192 × 640	0.099	0.708	4.372	0.175	0.900	0.967	0.984	
Ours-MonoViT		Tr	192 × 640	0.096	0.696	4.327	0.174	0.904	0.968	0.985	

Table 1. Performance comparison of self-supervised DE methods on Cityscapes and KITTI. ‘D’ column: Test dataset. †: multiple test frames. A checkmark on the ‘O.N.’: utilization of an additional object motion estimation network. ‘Tr’/‘Te’ in ‘Sem’ column: semantic information usage in training/testing. Metrics with \downarrow are the lower, the better, while metrics with \uparrow are the higher, the better.

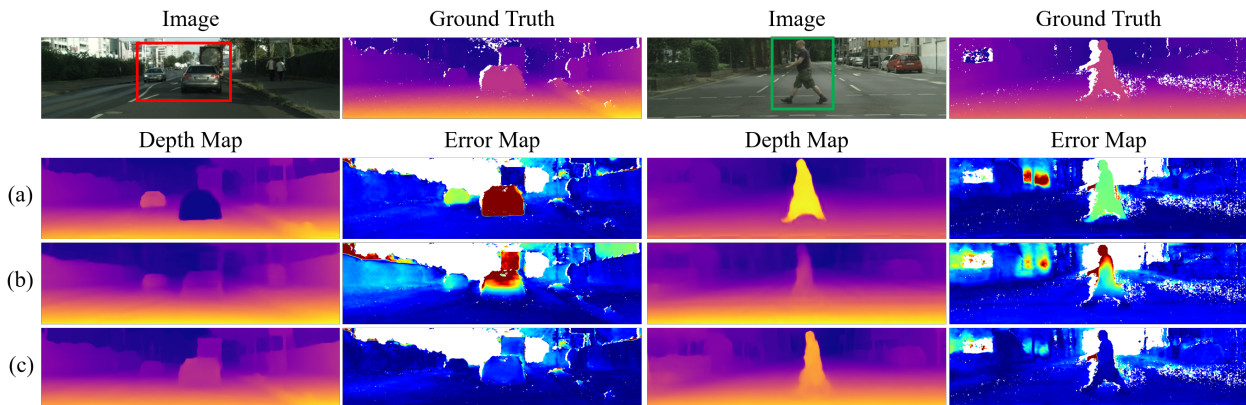


Figure 4. Examples of estimated depth maps and error maps in Cityscapes [6]. (a), (b), and (c) indicate MonoViT [54], DynamicDepth [9], and Ours-MonoViT, respectively. In the error maps, blue indicates small errors, and red indicates large errors.

4. Experiments

4.1. Datasets

Experiments are conducted on the widely adopted Cityscapes [6] and KITTI [38]. Cityscapes [6] contains a considerable amount of moving objects in the training dataset where self-supervised DE methods often suffer from the moving object problem. Following [51, 55], we use 69,731 training images (pre-processed with the scripts from [55]), and 1,525 testing images. For the DE performance on the regions of the objects in dynamic classes, we utilize masks provided by [9]. We also train the DE networks on the KITTI [38] dataset which contains 39,180 training and 4,424 validation images and evaluate on the Eigen split [8] which contains 697 images. For testing in both datasets, the estimated depth maps are normalized using median scaling within the depth range from 0m to 80m.

4.2. Implementation Details

Our coarse-to-fine training strategy is integrated with the monocular DE networks including Monodepth2 [13], HR-Depth [37] with ResNet-18 based encoder [21], CADepth [52] with ResNet-50 based encoder [21] and MonoViT [54] with MPViT-small based encoder, all pre-trained on ImageNet [28]. The DE network and the pose estimation network (with ResNet-18 [21] based encoder) are trained for 15 epochs in the coarse training stage and the DE network is refined for an additional 5 epochs in the fine training stage. Our overall training strategy including the cost volume computation (Eq. 12) with the fixed θ_{depth}^{1*} in the fine stage increased the total training time by about 11.1% (from 13.5h to 15h) on an NVIDIA RTX 3090 GPU for training Monodepth2 [13] in Cityscapes. The learning rates for ResNet-based models are set as 1×10^{-4} for the coarse stage and 1×10^{-5} for the fine stage, while Ours-MonoViT starts from 5×10^{-5} and decaying exponentially by a factor of 0.9 for 20 epochs. The overall pipeline is implemented in PyTorch [39] and trained with the Adam [26] optimizer with a batch size of 12. The loss weights α of the L1 loss (Eq. 3), β of the GDS-Loss (Eq. 10), and ρ of L_{REG} (Eq. 15) are empirically set to 0.85, 0.001 and 0.1, respectively. The masks of objects in dynamic classes for training are pre-computed by the Mask2Former [5] instance segmentation network with the backbone of Swin-S [36], pre-trained from COCO [34] dataset.

4.3. Cityscapes Results

Table 1 shows the performance comparison of self-supervised DE methods on Cityscapes and KITTI datasets. Especially on Cityscapes dataset which contains a substantial amount of moving objects, our training strategy significantly improves the DE performance of existing DE methods including Monodepth2 [13], HR-Depth [37], CADepth

Methods	Resolutions	<i>abs rel</i> ↓	<i>sq rel</i> ↓	<i>rmse</i> ↓	<i>log₁₀ rmse</i> ↓	<i>a¹</i> ↑
Monodepth2 [13]	128 × 416	0.159	1.937	6.363	0.201	0.816
MonoViT [54]	128 × 416	0.142	1.536	5.130	0.185	0.820
Ours-Monodepth2	128 × 416	0.136	1.238	4.791	0.176	0.864
DynamicDepth [9]	128 × 416	0.129	1.273	4.626	0.168	0.862
Ours-MonoViT	128 × 416	0.109	0.888	4.243	0.151	0.898
InstaDM [30]	256 × 832	0.139	1.698	5.760	0.181	0.859
Monodepth2 [13]	192 × 640	0.185	2.432	5.919	0.218	0.794
HR-Depth [37]	192 × 640	0.165	2.144	5.720	0.198	0.811
CADepth [52]	192 × 640	0.143	1.278	4.997	0.184	0.825
MonoViT [54]	192 × 640	0.149	1.508	5.340	0.190	0.817
Ours-Monodepth2	192 × 640	0.119	1.044	4.270	0.157	0.881
Ours-HR-Depth	192 × 640	0.117	1.015	4.297	0.154	0.883
Ours-CADepth	192 × 640	0.116	1.161	4.370	0.154	0.892
Ours-MonoViT	192 × 640	0.098	0.674	3.688	0.135	0.914

Table 2. Performance comparison of self-supervised DE methods for objects in dynamic classes on Cityscapes [6].

[52] and MonoViT [54]. At a resolution of 192×640 , the aforementioned methods are enhanced by incorporating our training strategy with large margins of 18.4%, 16.7%, 21.8%, and 17.0% in the *absrel* metric. Particularly, Ours-Monodepth2 outperforms other methods [37, 52, 54] that are trained without our training strategy. Also, Ours-CADepth and Ours-MonoViT outperform RM-Depth [22] without requiring additional object motion estimation networks. Notably, in the 128×416 resolution, our single-frame-based Ours-MonoViT outperforms the multi-frame-based SOTA DynamicDepth [9], which relies on segmentation masks in test time. In Table 2, we compare the DE performance on the regions of objects in dynamic classes. Our coarse-to-fine training strategy significantly enhances the DE performance of existing DE methods [13, 37, 52, 54]. Especially in a resolution of 128×416 , Ours-MonoViT outperforms DynamicDepth [9] in all metrics.

Also, in Fig. 4, MonoViT [54] and DynamicDepth [9] predict erroneous depth on a car (denoted in a red box) and a pedestrian (denoted in a green box) resulting in high errors in the error maps due to the lack of supervision for depth of moving objects. In contrast, Ours-MonoViT estimates accurate depth especially on the objects (the cars and the pedestrian) regardless of their moving speed and directions, while preserving details in depth maps.

4.4. KITTI Results

We further assess the efficacy of our coarse-to-fine training strategy on the frequently employed KITTI [38] dataset. As can be seen in Table 1, the performances of Monodepth2 [13], HR-Depth [37], CADepth [52] and MonoViT [54] are also improved with our training strategy integrated. Although KITTI has less amount of moving objects in the training scenes, our coarse-to-fine training strategy consistently boost the DE performance of existing DE methods. Also, Ours-MonoViT outperforms the current state-of-the-art multi-frame-based method, DynamicDepth [9].

Exp	L_{rep}^m	L_{GDS}		$abs\ rel \downarrow$	$sq\ rel \downarrow$	$rmse \downarrow$	$log_{rmse} \downarrow$	$a^1 \uparrow$
(a)			WIR	0.125	1.474	6.688	0.180	0.865
			DOR	0.185	2.432	5.919	0.218	0.794
(b)	✓		WIR	0.119	1.492	6.610	0.179	0.877
			DOR	0.160	2.143	5.503	0.203	0.813
(c)	✓	✓	WIR	0.104	1.097	6.034	0.160	0.895
			DOR	0.125	1.185	4.459	0.161	0.876

Table 3. Ablation study on coarse training stage by training Monodepth2 [13] on Cityscapes [6]. WIR and DOR indicate the performance measures over the whole image regions and over the regions of objects in dynamic classes *only*, respectively.

Exp.	Methods		$abs\ rel \downarrow$	$sq\ rel \downarrow$	$rmse \downarrow$	$log_{rmse} \downarrow$	$a^1 \uparrow$
(d)	No regularization	WIR	0.145	2.868	7.395	0.201	0.856
		DOR	0.303	10.533	8.839	0.281	0.764
(e)	Filtering [40]	WIR	0.124	1.620	6.739	0.187	0.868
		DOR	0.214	4.372	7.218	0.257	0.769
(f)	L_{REG}^0 (Eq. 11)	WIR	0.117	1.358	6.422	0.176	0.873
		DOR	0.181	2.594	6.047	0.218	0.791
(g)	L_{REG} (Eq. 14)	WIR	0.102	1.024	6.015	0.159	0.896
		DOR	0.119	1.044	4.270	0.157	0.881

Table 4. Ablation study on fine training stage.

5. Ablation Study

We demonstrate the ablation study to verify the effectiveness of our coarse-to-fine training strategy by incrementally applying the proposed contributions to Monodepth2 [13] and compare DE performance in Cityscapes [6] dataset.

Coarse Training Stage. In Table 3, from Exp-(a) to (c), we train the DE network for 15 epochs to show the effectiveness of our proposed masked reprojection loss L_{rep}^m and GDS-Loss L_{GDS} .

Exp-(a) - Original Monodepth2 [13]: As shown in Table 3, the automasking technique fails to guide the Monodepth2 [13] model to learn the accurate depth of moving objects since it is only capable of excluding the objects moving at the same speed and direction as the camera.

Exp-(b) - Monodepth2 + L_{rep}^m : Simply excluding the objects in dynamic classes from L_{rep} shows marginal improvement by avoiding inaccurate learning of depth from L_{rep} . However, it is not sufficient to guide precise learning of depth due to the lack of supervision in the object regions.

Exp-(c) - Monodepth2 + L_{rep}^m + L_{GDS} : Additional adoption of our GDS-Loss (Eq. 9) not only enhances the overall performance but also drastically improves the DE quality on the regions of objects in dynamic classes, since our GDS-Loss L_{GDS} provides precise depth supervision for the objects based on the ground contacting prior.

Fine Training Stage. In Table 4, from Exp-(d) to (g), we show the effectiveness of our regularization loss L_{REG} when the DE network θ_{depth}^2 , initially trained in Exp-(c), is

further refined with the reprojection loss L_{rep} to capture the detailed depth of objects in dynamic classes.

Exp-(d) - Training with no regularization: Although θ_{depth}^2 is initially trained in the coarse stage, the DE performance shows degradation due to the inaccurate guidance from L_{rep} in the moving object regions.

Exp-(e) - Usage of the filtering scheme in [40]: We mask out the pixels from L_{rep} , which are identified to be inconsistent in the 3D space by thresholding the difference between the depth of sequential frames predicted by θ_{depth}^{1*} . Although this approach is effective in filtering out noisy pixels from pseudo labels in [40], it struggles to effectively exclude moving objects from L_{rep} with thresholding so that θ_{depth}^2 suffers from inaccurate learning from L_{rep} unfiltered moving object regions.

Exp-(f) - Initial regularization loss L_{REG}^0 (Eq. 11): Although L_{REG}^0 moderately guides θ_{depth}^2 to predict consistent depth with θ_{depth}^{1*} , it struggles to effectively regularize the inaccurate supervision from L_{rep} , especially on the moving object regions.

Exp-(g) - Our regularization loss L_{REG} (Eq. 14) with λ_{cv} (Eq. 13): Since λ_{cv} effectively penalizes the moving object regions, θ_{depth}^2 successfully improves the DE performance with elaborating the depth details of the objects in dynamic classes. Moreover, the approach of weighting based on the depth discrepancy relative to the reliable θ_{depth}^{1*} is shown to be more stable than filtering pixels from L_{rep} that involves implicit uncertainty by discerning whether an object is moving or not.

6. Conclusion

We address a challenging moving object problem in self-supervised monocular depth estimation. For this, our proposed coarse-to-fine training strategy provides precise supervision on the depth of moving objects by firstly incorporating the *ground-contacting-prior-based self-supervision*, which is Ground-contacting-prior Disparity Smoothness Loss (GDS-Loss). The subsequent stage refines the DE network to capture the detailed depth of the objects in dynamic classes under well-defined constraints from our regularization loss with a cost-volume-based weighting factor. The extensive experimental results show that our training strategy is very well harmonized with existing DE methods to effectively handle such moving objects. Moreover, Ours-MonoViT yields the SOTA depth estimation performance on both KITTI and Cityscapes datasets.

Acknowledgements

This work was supported by IITP grant funded by the Korea government (MSIT) (No. RS2022-00144444, Deep Learning Based Visual Representational Learning and Rendering of Static and Dynamic Scenes).

References

- [1] Jinwoo Bae, Sungho Moon, and Sunghoon Im. Deep digging into the generalization of self-supervised monocular depth estimation. In *Proceedings of the AAAI conference on artificial intelligence*, pages 187–196, 2023. 2
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021. 1
- [3] 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*, pages 8001–8008, 2019. 2, 3, 6
- [4] 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/CVF International Conference on Computer Vision*, pages 7063–7072, 2019. 2
- [5] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 7, 3, 4
- [6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 6, 7, 8, 1, 2, 3, 4
- [7] Tom van Dijk and Guido de Croon. How do neural networks see depth in single images? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2183–2191, 2019. 2
- [8] 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. 7
- [9] Ziyue Feng, Liang Yang, Longlong Jing, Haiyan Wang, YingLi Tian, and Bing Li. Disentangling object motion and occlusion for unsupervised multi-frame monocular depth. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, pages 228–244. Springer, 2022. 1, 2, 3, 6, 7
- [10] 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. 1
- [11] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 740–756. Springer, 2016. 1, 2
- [12] 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. 1, 3
- [13] 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. 1, 2, 6, 7, 8, 3, 4, 5
- [14] Juan Luis Gonzalez and Munchurl Kim. Plade-net: Towards pixel-level accuracy for self-supervised single-view depth estimation with neural positional encoding and distilled matting loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6851–6860, 2021. 1, 2
- [15] Juan Luis GonzalezBello and Munchurl Kim. Forget about the lidar: Self-supervised depth estimators with med probability volumes. *Advances in Neural Information Processing Systems*, 33:12626–12637, 2020. 1, 2
- [16] 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/CVF International Conference on Computer Vision*, pages 8977–8986, 2019. 2, 6
- [17] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raveentos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 6, 4
- [18] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-guided representation learning for self-supervised monocular depth. In *International Conference on Learning Representations*, 2020. 1, 2
- [19] Vitor Guizilini, Rareş Ambruş, Dian Chen, Sergey Zakharov, and Adrien Gaidon. Multi-frame self-supervised depth with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 160–170, 2022. 2
- [20] Wencheng Han, Junbo Yin, Xiaogang Jin, Xiangdong Dai, and Jianbing Shen. Brnet: Exploring comprehensive features for monocular depth estimation. In *European Conference on Computer Vision*, pages 586–602. Springer, 2022. 2
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [22] Tak-Wai Hui. Rm-depth: Unsupervised learning of recurrent monocular depth in dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1675–1684, 2022. 1, 2, 6, 7
- [23] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Advances in neural information processing systems*, 28, 2015. 3
- [24] Adrian Johnston and Gustavo Carneiro. Self-supervised monocular trained depth estimation using self-attention and discrete disparity volume. In *Proceedings of the IEEE/CVF con-*

- ference on computer vision and pattern recognition*, pages 4756–4765, 2020. [2](#), [6](#)
- [25] Hyunyoung Jung, Eunhyeok Park, and Sungjoo Yoo. Fine-grained semantics-aware representation enhancement for self-supervised monocular depth estimation. *CoRR*, abs/2108.08829, 2021. [2](#), [6](#)
- [26] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. [7](#)
- [27] Marvin Klingner, Jan-Aike Termöhlen, Jonas Mikolajczyk, and Tim Fingscheidt. Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 582–600. Springer, 2020. [1](#), [2](#), [6](#)
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. [7](#)
- [29] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019. [1](#)
- [30] Seokju Lee, Sunghoon Im, Stephen Lin, and In So Kweon. Learning monocular depth in dynamic scenes via instance-aware projection consistency. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1863–1872, 2021. [1](#), [2](#), [3](#), [6](#), [7](#)
- [31] Seokju Lee, Francois Rameau, Fei Pan, and In So Kweon. Attentive and contrastive learning for joint depth and motion field estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4862–4871, 2021. [3](#), [6](#)
- [32] Hanhan Li, Ariel Gordon, Hang Zhao, Vincent Casser, and Anelia Angelova. Unsupervised monocular depth learning in dynamic scenes. In *Conference on Robot Learning*, pages 1908–1917. PMLR, 2021. [1](#), [2](#), [6](#)
- [33] Rui Li, Dong Gong, Wei Yin, Hao Chen, Yu Zhu, Kaixuan Wang, Xiaozhi Chen, Jinqiu Sun, and Yanning Zhang. 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*, pages 21539–21548, 2023. [1](#)
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. [7](#)
- [35] Yuxuan Liu, Yuan Yixuan, and Ming Liu. Ground-aware monocular 3d object detection for autonomous driving. *IEEE Robotics and Automation Letters*, 6(2):919–926, 2021. [1](#)
- [36] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. [7](#)
- [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. *CoRR abs/2012.07356*, 2020. [2](#), [6](#), [7](#), [3](#), [4](#), [5](#)
- [38] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. [7](#), [1](#), [2](#), [3](#), [4](#)
- [39] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [7](#)
- [40] Andra Petrovai and Sergiu Nedevschi. Exploiting pseudo labels in a self-supervised learning framework for improved monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1578–1588, 2022. [2](#), [6](#), [8](#)
- [41] Andrea Pilzer, Stephane Lathuiliere, Nicu Sebe, and Elisa Ricci. Refine and distill: Exploiting cycle-inconsistency and knowledge distillation for unsupervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9768–9777, 2019. [2](#)
- [42] Matteo Poggi, Fabio Tosi, and Stefano Mattoccia. Learning monocular depth estimation with unsupervised trinocular assumptions. In *2018 International conference on 3d vision (3DV)*, pages 324–333. IEEE, 2018. [2](#)
- [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. [2](#)
- [44] Zequn Qin and Xi Li. Monoground: Detecting monocular 3d objects from the ground. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3793–3802, 2022. [1](#)
- [45] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12179–12188, 2021. [1](#)
- [46] 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. [2](#)
- [47] Kunal Swami, Amrit Muduli, Uttam Gurram, and Pankaj Bajpai. Do what you can, with what you have: Scale-aware and high quality monocular depth estimation without real world labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 988–997, 2022. [2](#)
- [48] Ruoyu Wang, Zehao Yu, and Shenghua Gao. Planedepth: Self-supervised depth estimation via orthogonal planes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21425–21434, 2023. [1](#)
- [49] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4): 600–612, 2004. [3](#)

- [50] Jamie Watson, Michael Firman, Gabriel J Brostow, and Daniyar Turmukhambetov. Self-supervised monocular depth hints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2162–2171, 2019. [1](#), [2](#)
- [51] Jamie Watson, Oisin Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1164–1174, 2021. [1](#), [2](#), [5](#), [7](#)
- [52] Jiaxing Yan, Hong Zhao, Penghui Bu, and YuSheng Jin. Channel-wise attention-based network for self-supervised monocular depth estimation. In *2021 International Conference on 3D vision (3DV)*, pages 464–473. IEEE, 2021. [2](#), [6](#), [7](#), [3](#), [4](#), [5](#)
- [53] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 340–349, 2018. [2](#)
- [54] Chaoqiang Zhao, Youmin Zhang, Matteo Poggi, Fabio Tosi, Xianda Guo, Zheng Zhu, Guan Huang, Yang Tang, and Stefano Mattoccia. Monovit: Self-supervised monocular depth estimation with a vision transformer. In *International Conference on 3D Vision*, 2022. [1](#), [2](#), [5](#), [6](#), [7](#), [3](#), [4](#)
- [55] 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. [2](#), [7](#)
- [56] Yunsong Zhou, Quan Liu, Hongzi Zhu, Yunzhe Li, Shan Chang, and Minyi Guo. Mogde: Boosting mobile monocular 3d object detection with ground depth estimation. *Advances in Neural Information Processing Systems*, 35:2033–2045, 2022. [1](#)
- [57] Zhengming Zhou and Qiulei Dong. Self-distilled feature aggregation for self-supervised monocular depth estimation. In *European Conference on Computer Vision*, pages 709–726. Springer, 2022. [2](#)