

Camera Height Doesn't Change: Unsupervised Training for Metric Monocular Road-Scene Depth Estimation

Genki Kinoshita[✉] and Ko Nishino[✉]

Graduate School of Informatics, Kyoto University, Kyoto, Japan
gkinoshita@vision.ist.i.kyoto-u.ac.jp, kon@i.kyoto-u.ac.jp
<https://vision.ist.i.kyoto-u.ac.jp/>

Abstract. In this paper, we introduce a novel training method for making any monocular depth network learn absolute scale and estimate metric road-scene depth just from regular training data, *i.e.*, driving videos. We refer to this training framework as FUMET. The key idea is to leverage cars found on the road as sources of scale supervision and to incorporate them in network training robustly. FUMET detects and estimates the sizes of cars in a frame and aggregates scale information extracted from them into an estimate of the camera height whose consistency across the entire video sequence is enforced as scale supervision. This realizes robust unsupervised training of any, otherwise scale-oblivious, monocular depth network so that they become not only scale-aware but also metric-accurate without the need for auxiliary sensors and extra supervision. Extensive experiments on the KITTI and the Cityscapes datasets show the effectiveness of FUMET, which achieves state-of-the-art accuracy. We also show that FUMET enables training on mixed datasets of different camera heights, which leads to larger-scale training and better generalization. Metric depth reconstruction is essential in any road-scene visual modeling, and FUMET democratizes its deployment by establishing the means to convert any model into a metric depth estimator.

Keywords: Unsupervised Monocular Depth Estimation · Metric Scale · Camera Height Invariance · Object Size Prior

1 Introduction

Monocular depth estimation is essential for autonomous driving and advanced driver assistance systems (ADAS). It underlies many of their key perceptual tasks including 3D object detection [7, 33, 39, 49], motion planning [31, 38, 51], and road segmentation [26]. Although supervised methods have achieved impressive accuracy, they suffer from costly ground-truth data collection. This limits their generalizability [34] and also makes continuous learning difficult.

Many recent methods [5, 15, 18, 25, 52] leverage self-supervision to scale up training. Most of these methods, however, suffer from scale ambiguity which renders them difficult to use off-the-shelf. To resolve the scale ambiguity and

recover in metric scale, the actual scene scale needs to be supervised in some form. Past works have leveraged velocity [16], GPS [3], IMU [41, 53], or camera height [36, 45] for this. These methods require auxiliary sensors to measure these extra quantities in addition to the monocular RGB camera, which significantly limits the ability to scale the size of training data. Arbitrary in-the-wild driving videos cannot be used as they lack such measurements. These self-supervised methods often still require weak scale supervision and cannot be made fully self-supervised as the scale changes for each sequence captured with a different camera.

Road scenes actually have plenty of untapped sources of metric scale. Cars found on the road are one of them as they are rigid objects whose actual sizes do not change and are unique to each make and model. The road is also exactly the region where we need accurate metric depth the most. If we can leverage the sizes of cars to bridge the 3D to 2D projection, we may train a mono-depth network to become metric-accurate. Simply employing a prior on the car size, however, would be too brittle since the metric supervision will be as ambiguous as the accuracy of that prior. The perspective projection and the variations in car makes, years, and colors, lead to a large variance in the size estimates, which will directly affect the mono-depth metric accuracy. In addition, there are no reliable object size priors except for a few introduced as probability densities for other tasks [20, 54, 57] whose parameters are not shared.

How can we train an arbitrary mono-depth estimator to become metric-accurate by robustly leveraging prior knowledge about the 3D sizes of cars found in driving videos? We derive **fully unsupervised metric depth trainer (FUMET)** to answer this question. Our key idea is to aggregate car size cues in the frames and the sequence into a single physical measure, namely the camera height. Regardless of the scene, the camera capturing it is the same for the same sequence. This means that, by transforming the car size cues found in different frames into a camera height estimate, the very fact that it should not change can be used as a supervision of invariance to achieve metric depth reconstruction. We can formulate this as an optimization of the camera height across frames and training epochs. This leads to gradual metric-aware learning that is stable and accurate, which also makes the depth estimates consistent across the sequence.

We introduce a **learned size prior (LSP)** to estimate the size of each car found in a frame. This prior gives us the vehicle dimensions from its appearance. It is trained with various augmentations on a large-scale dataset without the need for any manual annotation. By comparing these dimension estimates with those computed from the depth estimates, we obtain a per-frame scale factor. The multiplication of this scale estimate with the camera height estimate derived from the estimated depth becomes the camera height estimate of the frame. These estimated camera heights across all frames in a sequence are then consolidated by using their median value which is then used as scale supervision with its weighted moving average at the end of each training epoch.

We evaluate the effectiveness of FUMET with extensive experiments using the KITTI and the Cityscapes datasets and compare its accuracy with a number

of weakly-supervised methods. The results clearly show that FUMET succeeds in making monocular depth estimation methods learn metric scales and achieves the state-of-the-art accuracy without any direct depth and scale supervision. We also demonstrate its capability to train a model on mixed datasets collected with different cameras, *i.e.*, camera heights. This enables larger-scale training and leads to higher generalization. We believe FUMET can serve as a fundamental building block for monocular depth estimation in the wild.

2 Related work

2.1 Self-Supervision

Although, as in other tasks, supervised methods achieve the highest accuracy in monocular depth estimation (MDE) [1, 23, 27, 29, 50], they obviously require ground-truth data usually obtained with a LiDAR sensor. Self-supervised MDE methods remove the need for such hardware costs for training. Garg *et al.* [11] proposed the first approach to self-supervised MDE by predicting per-pixel disparities between a stereo pair. Godard *et al.* [14] extend this idea by introducing left-right consistency and greatly improve the accuracy. These methods require calibrated stereo data for training which limits their applicability. Zhou *et al.* [56] introduced a pioneering method for training solely from monocular videos by incorporating a network that estimates the relative camera poses between successive frames. Based on this idea, many MDE methods that can be trained solely with monocular videos have been proposed [5, 15, 18, 25, 52, 55].

Despite the popularity and improvements of self-supervised MDE, directly deploying those models in autonomous driving or ADAS systems is unrealistic as they only produce relative depth up to an unknown scale. Our FUMET overcomes this issue without the need for any external supervision.

2.2 Weak Supervision

Several works have tackled scale-aware MDE by adopting weak scale supervision through velocity [16], GPS [3], IMU [41], or camera height [36, 45] measurements. Zhang *et al.* [53] use the combination of IMU, velocity, and gravity to achieve the highest accuracy among these methods. These methods fundamentally rely on such auxiliary measurements in addition to regular RGB images. This makes it impossible to take advantage of the many road-scene videos on the Internet (*e.g.*, YouTube) or casually installed dashcam videos for training. In contrast, FUMET relies on a universal car size prior which is independent of the data collection environments and does not require any such weak supervision; it can be trained solely from monocular video sequences.

2.3 Object Size Priors

Several works for other tasks leverage priors on object sizes including single view metrology [57], object insertion [20, 54], and visual odometry [10, 35]. These

methods model the priors as normal distributions or Gaussian mixture models, and the parameters of these distributions are predetermined. When there is a gap between those values and the true ones, which is usually the case in the real world, the difference directly surfaces in erroneous scale estimates.

Instead of these fixed priors, we introduce a learning-based size prior which we refer to as LSP. This estimates vehicle dimensions based on its appearance, thus can avoid the aforementioned problem. Furthermore, it is trained on a large-scale dataset with various augmentations, so that it can robustly estimate vehicle dimensions from any image.

3 Preliminaries

Our goal is to train a metric-accurate MDE model θ_{depth} with only self-supervision. This self-supervision is comprised of two parts: one designed for learning metric scale (Sec. 4) and the other dedicated to learning depth geometry. For the latter, we follow Zhou *et al.* [56] and minimize the reconstruction loss between the current frame I_t and the ones synthesized from temporally adjacent frames I_s , ($s = t \pm 1$). We employ a pose network θ_{pose} in addition to θ_{depth} to estimate the relative camera pose $\mathbf{T}_{t \rightarrow s} \in \text{SE}(3)$ and the current dense depth map D_t , respectively,

$$D_t = \theta_{\text{depth}}(I_t), \quad \mathbf{T}_{t \rightarrow s} = \theta_{\text{pose}}(I_t, I_s). \quad (1)$$

We then synthesize the image from the current viewpoint

$$I_{s \rightarrow t} = I_s \langle \text{proj}(D_t, \mathbf{T}_{t \rightarrow s}, \mathbf{K}) \rangle, \quad (2)$$

where $\langle \cdot \rangle$ is a bilinear sampling operator, $\text{proj}(\cdot)$ outputs the 2D coordinates of D_t when projected into the viewpoint of I_s , and \mathbf{K} is the camera intrinsics. Following Godard *et al.* [15], we use an affine combination of SSIM [42] and L1 loss as the photometric error pe and define the reconstruction loss \mathcal{L}_{rec} as the photometric error between I_t and $I_{s \rightarrow t}$ for the pixels with minimum error across source frames

$$\mathcal{L}_{\text{rec}} = \min_s pe(I_t, I_{s \rightarrow t}), \quad (3)$$

$$pe(I_t, I_{s \rightarrow t}) = \frac{\lambda}{2}(1 - \text{SSIM}(I_t, I_{s \rightarrow t})) + (1 - \lambda)|I_t - I_{s \rightarrow t}|, \quad (4)$$

where $\lambda = 0.85$. We also employ an edge-aware smoothness loss [14]

$$\mathcal{L}_{\text{sm}} = |\partial_x d_t^*| e^{-|\partial_x I_t|} + |\partial_y d_t^*| e^{-|\partial_y I_t|}, \quad (5)$$

where $d_t^* = d_t / \bar{d}_t$ is the mean-normalized inverse depth [37] to prevent the estimated depth from shrinking. Furthermore, when computing \mathcal{L}_{rec} we use auto-masking [15] to mask out stationary pixels which generate incorrect supervisory signals.




Fig. 1: Overview of FUMET. At each training step n , an *unscaled* camera height $H_{\text{cam}}'^{n, \tau}$ is computed differentiably from the estimated depth. The previous epoch $\tau-1$ provides supervision with a scaled camera height $H_{\text{cam}}^{* \tau-1}$. To obtain this scaled camera height supervision, Silhouette Projector first computes the object silhouette heights H'_{obj} from the depth map. By comparing H'_{obj} and the estimated one H_{obj} with LSP, per-frame scale factor s is determined and we obtain the scaled camera height $H_{\text{cam}}^{n, \tau} = s \cdot H_{\text{cam}}'^{n, \tau}$. At the end of each epoch, H_{cam}^* is optimized across a series of consecutive frames and updated with the weighted moving average.

4 FUMET

Fig. 1 depicts the overall architecture of FUMET. Please see the caption for an overview of the framework. Let “'” denote “not metrically scaled”.

4.1 Scale-Aware Self-Supervised Learning

In FUMET, we derive a camera height pseudo-supervision from depth estimates. To obtain the per-frame camera height, we first compute the normal vector for each pixel p_i . We consider the 8-neighborhood of p_i and form 8 pairs $\mathcal{N}(p_i) = \{(p_{j_0}^i, p_{j_1}^i)\}_{j=1}^8$. Each pair $(p_{j_0}^i, p_{j_1}^i)$ consists of two pixels whose difference vectors w.r.t. p_i are perpendicular to each other and are sorted counterclockwise. We reproject the paired pixels onto 3D space and compute the cross-product of the two difference vectors w.r.t. the reprojected 3D point of p_i . By summing and normalizing all the cross-product vectors of each pair, we obtain the normal vector $\hat{\mathbf{n}}(p_i)$

$$\mathbf{n}(p_i) = \sum_{\mathcal{N}(p_i)} (\phi(p_{j_0}^i) - \phi(p_i)) \times (\phi(p_{j_1}^i) - \phi(p_i)), \quad \hat{\mathbf{n}}(p_i) = \frac{\mathbf{n}(p_i)}{\|\mathbf{n}(p_i)\|}, \quad (6)$$

where $\phi(\cdot)$ denotes perspective reprojection.

For pixels in the road region, the negated inner product between the normal vector and the reprojected 3D point corresponds to the camera height from the ground plane

$$H'_{\text{cam}}(p_i) = -\phi(p_i) \cdot \hat{\mathbf{n}}(p_i). \quad (7)$$

We represent the median of $H'_{\text{cam}}(p_i)$ as the camera height of the frame

$$\tilde{H}'_{\text{cam}} = \text{med}(\{H'_{\text{cam}}(p_i) \mid p_i \in M_r\}), \quad (8)$$

where p_i is in the road region mask M_r obtained with an off-the-shelf semantic segmentation model [22] before training. This camera height \tilde{H}'_{cam} estimate is *unscaled* and needs to be upgraded to a real-scaled value to impose the invariance loss across frames. By using the Silhouette Projector, we aggregate the scale information from the object height prior into a per-frame scale factor s (see Sec. 4.2) and obtain the *scaled* camera height as $H_{\text{cam}} = s \cdot \tilde{H}'_{\text{cam}}$.

We leverage the fact that the camera height does not change in a sequence and formulate the invariance into an optimization across frames and training epochs. Thanks to this optimization, we acquire an accurate and stable supervision, which is also consistent across the scene. At the end of the τ -th epoch, we use the median of $\{H_{\text{cam}}^{n,\tau}\}_{n=1}^N$ for all frames N as the representative camera height H_{cam}^τ at the τ -th epoch. Then, we update the *scaled* camera height pseudo-invariance H_{cam}^* with its weighted moving average

$$H_{\text{cam}}^{*\tau} \leftarrow \left\{ \frac{\tau(\tau-1)}{2} \cdot H_{\text{cam}}^{*\tau-1} + \tau \cdot H_{\text{cam}}^\tau \right\} / \frac{\tau(\tau+1)}{2}. \quad (9)$$

When training on a dataset containing images captured at multiple camera heights, we create pseudo-labels corresponding to each camera height and optimize them individually.

To make the network learn the absolute scale and directly produce the *scaled* depth map, we enforce a loss on the camera height H'_{cam} using the pseudo-invariance $H_{\text{cam}}^{*\tau-1}$ computed in the previous epoch

$$\mathcal{L}_{\text{cam}} = \frac{1}{|M_r|} \sum_i M_r(p_i) \cdot |H'_{\text{cam}}(p_i) - H_{\text{cam}}^{*\tau-1}|. \quad (10)$$

It is worth noting that at inference, the network does not assume camera height invariance in the input sequence; it is only assumed during training. The computational cost at inference is exactly the same as the original MDE model trained without FUMET.

4.2 Silhouette Projector

In the process of creating a *scaled* camera height supervision, we need to know the per-frame scale factor. For this, we can consider leveraging an object of known size. This, however, would necessitate extraction of the exact 3D object region which is non-trivial as objects are often occluded or cropped and oriented in various poses.

We introduce Silhouette Projector to robustly estimate this per-frame scale factor. Silhouette Projector leverages two simple facts of road scenes. First, the height of the object silhouette projected on the plane perpendicular to the ground plane does not change regardless of the object pose. Second, we can compute the silhouette height by measuring the distance from the top of the silhouette to the ground plane even if they are partly occluded or truncated, as long as the top is visible and the object is on the ground plane.

We first reconstruct an object point cloud with the estimated depth map through perspective reprojection. The object region M_{obj} in the image can be obtained with an off-the-shelf instance segmentation model [22] before training. We then orthographically project these 3D points onto a plane P^\perp perpendicular to the ground plane, *i.e.*, an arbitrary plane parallel to the road normal vector $\tilde{\mathbf{n}}$. We compute $\tilde{\mathbf{n}}$ as the median of the normal vectors $\hat{\mathbf{n}}(p_i)$ (see Eq. (6)) in the road region M_r . We define the farthest distance from the 2D point on P^\perp to the ground plane as the silhouette height. This value is *unscaled* because we derived this from the estimated *unscaled* depth map. The ratio between a height prior and its silhouette height corresponds to the scale factor of interest. When there are multiple objects in an image, we use the median of the scale factors each corresponding to each object as the scale factor of the frame for computational simplicity and robustness.

Outlier Filtering. Some objects cannot be handled well with Silhouette Projector for various reasons including occluded upper parts, not on the ground plane, inaccurate segmented region, and misclassified object category. We introduce an algorithm to detect these as outliers based on geometric plausibility. We first calculate the horizon \mathbf{l} with the road normal vector $\tilde{\mathbf{n}}$ and the camera intrinsics \mathbf{K} [17]

$$\mathbf{l} = \mathbf{K}^{-\top} \tilde{\mathbf{n}}. \quad (11)$$

Based on the approximations in [20], we derive the approximate metric-scale object height $\hat{H}_{\text{obj}} \approx \frac{h_{\text{obj}}}{h_{\text{cam}}} H_{\text{cam}}^{*\tau-1}$, where h_{cam} is the farthest distance from the pixel in the object instance mask to \mathbf{l} , h_{obj} is the farthest distance between two points on the object mask on which tangent lines parallel to \mathbf{l} lie, and $H_{\text{cam}}^{*\tau-1}$ is the pseudo camera height at the τ -th epoch derived in Eq. (9). By comparing the relative gap between \hat{H}_{obj} and the object height prior H_{obj} , we determine the outliers as

$$\left\{ k \left| \frac{|H_{\text{obj}}^k - \hat{H}_{\text{obj}}^k|}{H_{\text{obj}}^k} > \mathcal{T} \right. \right\}. \quad (12)$$

In all the experiments, we set the threshold \mathcal{T} to 0.2.

4.3 Training Losses

Auxiliary Rough Geometric Loss. We experimentally observe that estimated object depths, especially those distant from the camera, often erroneously lean towards the road surface because of the inaccurate supervisory signal from the reconstruction loss in low-texture regions. This tendency makes silhouette heights inaccurate. To mitigate this issue, we introduce an auxiliary rough geometric loss \mathcal{L}_{aux} , which enforces object point clouds to be within a close region. For each inlier object k ($k = 1 \dots K$), we define \mathcal{L}_{aux} as the gap between the estimated depth map D in the object region M_{obj} and the approximate object

depth $D_{\text{aprx}} = \frac{H_{\text{obj}}}{h_{\text{obj}}} f_y$ where f_y is the focal length

$$\mathcal{L}_{\text{aux}} = \frac{1}{K} \sum_k \frac{1}{|M_{\text{obj}}^k|} \sum_i M_{\text{obj}}^k(p_i) \cdot \left| D(p_i) - D_{\text{aprx}}^k \right|. \quad (13)$$

As a by-product, this loss accelerates learning the metric scale.

The \mathcal{L}_{aux} is potentially inaccurate since D_{aprx} assumes that each object can be represented by a plane parallel to the image plane, and we observe that it degrades the depth accuracy in the late stage of training. On the other hand, in the early stage, estimated depth is unreliable, so heavily relying on the camera height loss \mathcal{L}_{cam} makes the training unstable. For this, we gradually decrease the weight of \mathcal{L}_{aux} and increase the one for \mathcal{L}_{cam} as the training proceeds in a logarithmic way. To ensure training stability, we stop changing these weights after the τ_{mid} -th epoch

$$\lambda_{\text{aux}}(\tau) = \begin{cases} \frac{-\log(\tau+1)}{\log(\tau_{\text{mid}}+1)} & (\tau \leq \tau_{\text{mid}}) \\ \epsilon & (\tau > \tau_{\text{mid}}) \end{cases}, \quad \lambda_{\text{cam}}(\tau) = \begin{cases} \frac{\log(\tau+1)}{\log(\tau_{\text{mid}}+1)} & (\tau \leq \tau_{\text{mid}}) \\ 1.0 & (\tau > \tau_{\text{mid}}) \end{cases}, \quad (14)$$

where $\epsilon = 0.005$.

The total loss is $\mathcal{L} = \alpha \lambda_{\text{cam}} \mathcal{L}_{\text{cam}} + \beta \lambda_{\text{aux}} \mathcal{L}_{\text{aux}} + \mathcal{L}_{\text{sm}} + \mathcal{L}_{\text{rec}}$ where $\alpha = 0.01$ and $\beta = 1.0$.

5 LSP

Our FUMET extracts metric scale from a prior on object heights through the camera height loss \mathcal{L}_{cam} and the auxiliary rough geometric loss \mathcal{L}_{aux} . To ensure that FUMET learns the scale regardless of the particular dataset used for training, a universal prior of object heights, independent of the environment, becomes essential. Cars are objects whose dimensions are explicitly and accurately known. Better yet, they are abundantly located on the road, the key semantic region whose metric depth is essential for many downstream tasks. Some methods for other tasks have modeled the size prior as a probability distribution [20, 57]. The gap between its mean and the true expectation for each environment, however, is directly reflected in the uncertainty of the scale in the supervision. Wang *et al.* [40] also report that there is a gap in car dimensions among datasets. These facts suggest that modeling the object size prior as a single probability distribution is problematic, especially for man-made objects like cars.

We humans not only possess rough prior knowledge about the vehicle size but can also estimate it more accurately by extracting instance-specific information such as car models from its appearance. Motivated by this intuition, we introduce a learned universal vehicle size prior termed LSP. Given a masked vehicle image, LSP estimates its dimensions.

For training, we use a dataset consisting of a large number of vehicle images and their information such as 3D dimensions. This dataset can be easily collected with web scraping and does not require manual annotations. To enhance

the robustness to various input images, in addition to common color jittering and blur augmentation, we perform augmentations to simulate occlusion and truncation. For occlusion augmentation, we mask out random regions with the same shape as a vehicle or with a beam of random width. For truncation augmentation, we translate the vehicle mask in a random direction. We visualize these in the supplementary material. The abundance of training images and diverse augmentations enable LSP to estimate vehicle dimensions robustly across various datasets.

Although FUMET requires only a car height prior, we construct LSP to estimate the width and length as well as height. This enhances it to grasp an overall car size and improves the prediction accuracy of car heights.

6 Experiments

6.1 FUMET

We evaluate the accuracy of depth estimators trained on our framework and demonstrate the effectiveness of FUMET as a novel unsupervised metric-scale training method through depth prediction accuracy, its applicability to any network architecture, and its ability to leverage any dataset for training.

Implementation Details. We implement FUMET with PyTorch [28]. We employ Monodepth2 [15] with ResNet50 encoder pre-trained on ImageNet [30] as the default depth estimator which is one of the most basic architectures in self-supervised MDE. For all experiments unless mentioned otherwise, we train the model for 50 epochs with a batch size of 8 and set τ_{mid} to 20. We use Adam optimizer [24] with an initial learning rate of 5.0×10^{-5} decayed by half every 15 epochs. During training, we perform the following augmentations with 50% chance in random order: horizontal flips, brightness adjustment (± 0.2), saturation adjustment (± 0.2), contrast adjustment (± 0.2), and hue jitter (± 0.1). For quantitative evaluation, we follow [9] and compute the seven standard metrics: AbsRel, SqRel, RMSE, RMSE_{\log} , $\delta < 1.25$, $\delta < 1.25^2$, and $\delta < 1.25^3$.

KITTI. Our FUMET is independent of network architectures so we experiment on several models with FUMET as well as Monodepth2: Lite-Mono [52] with ResNet18, HR-Depth [25] with ResNet18, and VADepth [45]. We train these models on the KITTI raw dataset [12], which contains 39,810 images for training and 4,424 for validation, and evaluate it with the KITTI Eigen test split [8] (697 images). For training and evaluation, all images are resized to 640×192 .

We show the qualitative results in Fig. 2. The quantitative results shown in Tab. 1 demonstrate that FUMET enables the various architectures to learn metric scale properly. Even when employing the simplest network architecture, *i.e.*, Monodepth2, it outperforms weakly-supervised methods that rely on ground-truth scale labels for training which is not necessary for our method. Similar to VADepth [45], FUMET uses the camera height invariance as supervision though

Table 1: Quantitative results on the KITTI Eigen test split [8] (640×192 resolution). In the Supervision column, V, G, and CamH represent velocity, gravity, and camera height, respectively. The Scaling column represents whether to perform the median scaling [56]. We denote the best results in **bold** and the second best ones with an underline for each block. Note that we train VADepth [45] from scratch using ground-truth camera height (1.65m) provided by KITTI, unlike its original paper. FUMET outperforms weakly-supervised methods regardless of the network architectures and without the need for ground truth supervision. In the rescaled results, every model trained with FUMET achieves better results than its original one.

Method	Supervision	Scaling	Error (↓)				Accuracy (↑)		
			AbsRel	SqRel	RMSE	RMSE _{log}	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
G2S [3]	GPS	-	0.109	0.860	4.855	0.198	0.865	0.954	0.980
PackNet-SfM [16]	V	-	0.111	0.829	4.788	0.199	0.864	0.954	0.980
Wagstaff <i>et al.</i> [36]	CamH	-	0.123	0.996	5.253	0.213	0.840	0.947	0.978
VADepth [45]	CamH	-	0.120	0.975	4.971	0.203	0.867	0.956	0.979
DynaDepth [53]	IMU+V+G	-	<u>0.109</u>	<u>0.787</u>	4.705	0.195	0.869	0.958	<u>0.981</u>
Ours	-	-	0.108	0.785	4.736	0.195	<u>0.871</u>	0.958	0.981
Lite-Mono [52]+ Ours	-	-	0.112	0.798	<u>4.692</u>	0.194	<u>0.871</u>	<u>0.959</u>	<u>0.981</u>
HR-Depth [25]+ Ours	-	-	<u>0.109</u>	0.794	4.734	<u>0.193</u>	0.869	<u>0.959</u>	0.982
VADepth [45]+ Ours	-	-	0.108	0.809	4.572	0.185	0.883	0.963	0.982
G2S [3]	GPS	✓	0.112	0.894	4.852	0.192	0.877	0.958	0.981
PackNet-SfM [16]	V	✓	0.111	0.785	4.601	0.189	0.878	0.960	0.982
Wagstaff <i>et al.</i> [36]	CamH	✓	0.117	0.952	4.989	0.197	0.867	0.957	0.981
VADepth [45]	CamH	✓	0.110	0.977	4.872	<u>0.187</u>	0.889	<u>0.962</u>	0.981
DynaDepth [53]	IMU+V+G	✓	<u>0.108</u>	<u>0.761</u>	4.608	<u>0.187</u>	0.883	<u>0.962</u>	0.982
Ours	-	✓	0.106	0.759	<u>4.602</u>	0.184	<u>0.887</u>	0.963	0.983
Monodepth2 [15]	-	✓	0.115	0.903	4.863	0.193	0.877	0.959	0.981
+ Ours	-	✓	0.106	0.759	4.602	0.184	0.887	0.963	0.983
HR-Depth [25]	-	✓	0.109	0.792	4.632	0.185	0.884	0.962	0.983
+ Ours	-	✓	0.108	0.759	4.584	0.184	0.885	0.963	0.983
Lite-Mono [52]	-	✓	0.110	0.802	4.671	0.186	0.879	0.961	0.982
+ Ours	-	✓	0.107	0.775	4.564	0.183	0.888	0.963	0.983
VADepth [45]	-	✓	0.104	0.774	4.552	0.181	0.892	0.965	0.982
+ Ours	-	✓	0.102	0.760	4.473	0.178	0.897	0.966	0.983

VADepth requires the ground-truth camera height and employs a complex network architecture. FUMET with Monodepth2 (row 6 in Tab. 1) achieves higher accuracy. This indicates that measuring camera height precisely is difficult and using unreliable ground-truth camera height causes low accuracy in the scale of estimated depth maps.

In the results showing the accuracy of the estimated depth map scaled with ground-truth scale factors via median scaling [56], it is evident that every model trained with FUMET attains superior results to the one trained without FUMET. This shows that our training framework not only enables the models to learn metric scale but also enhances the geometric accuracy. This improvement owes to the camera height loss \mathcal{L}_{cam} and \mathcal{L}_{aux} , which makes the model respect road regions and object regions, respectively.

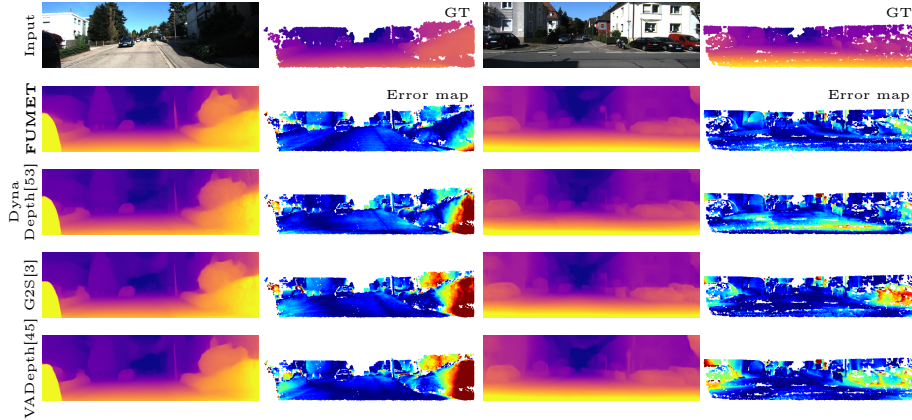


Fig. 2: Qualitative comparison on KITTI. In error maps, the larger depth errors are represented in red, smaller ones are in blue. The model trained with our FUMET predicts more accurate depth maps compared to the weakly-supervised methods.

FUMET successfully enables a model to learn metric scale, even though the training dataset of KITTI contains many frames with no or few detected cars. We detail this in the supplementary material.

Cityscapes. To evaluate the ability of FUMET to leverage any dataset for training, we train and evaluate it on the Cityscapes dataset [6]. We follow [43, 46, 48, 56] and collect 69,731 images for training and evaluate on the 1,525 test images. In this dataset, there are multiple sets of camera intrinsics. Since this strongly affects the metric depth estimation [47], we first align the focal lengths of all images to 587.5 by resizing them, and crop them to 512×192 resolution. For evaluation, we follow [2, 43] and further crop the estimated image into 416×128 . In order to align the number of iterations on KITTI, we train for 29 epochs with the same learning rate decreased by half every 8 epochs, and τ_{mid} is set to 12.

Tab. 2 shows the quantitative results and FUMET outperforms weakly-supervised methods by wide margins. These methods assume that highly accurate sensor data is available and fail to train scale-aware MDE models on datasets consisting of unreliable sensor data such as Cityscapes. On the other hand, FUMET does not require any sensor measurement other than RGB video, so it can robustly leverage any driving videos for training.

Mixed datasets. To demonstrate that FUMET can train a model on a mixed set of datasets to attain strong generalizability, we train with 235,341 images collected from the Argoverse2 [44], Lyft [21], A2D2 [13], and DDAD [16] dataset and evaluate it on the KITTI Eigen test split [8]. Both images for training and

Table 2: Quantitative results on Cityscapes [6]. We train all the models from scratch with the same supervisions as the Supervision column in Tab. 1. FUMET significantly outperforms weakly-supervised methods.

Method	AbsRel	SqRel	RMSE	$ \delta < 1.25$
G2S [3]	4.156	276.16	58.89	0.046
PackNet-SfM [16]	0.504	6.639	14.90	0.029
VADepth [45]	0.363	7.115	11.95	0.295
FUMET	0.125	1.288	6.359	0.858

Table 3: FUMET can be trained on mixed datasets (Argoverse2, Lyft, A2D2, and DDAD) with various camera settings, and attains high generalizability. The results are evaluated on the KITTI Eigen split [8] (832×512 resolution).

Training dataset	AbsRel	SqRel	RMSE	$ \delta < 1.25$
KITTI	0.103	0.675	4.708	0.903
Mixed	0.113	0.916	5.009	0.883
Mixed + KITTI	0.082	0.611	4.307	0.923

evaluation are resized and cropped to 832×512 resolution. The training images are the sets of video sequences and each sequence is captured with its own unique camera setting. We individually optimize each camera height for each sequence during training. We detail the dataset composition and training settings in the supplementary material.

As shown in Tab. 3, the model trained with FUMET on the large-scale mixed dataset achieves comparable accuracy to the one trained on KITTI despite the domain gap and attains high generalizability. The model trained on both the mixed dataset and KITTI outperforms the model trained only on KITTI. This suggests that training on various images enhances accuracy which is the strength of fully self-supervised learning methods.

We also show that FUMET can leverage in-the-wild videos for training by using camera intrinsics recovered with SLAM, *e.g.*, COLMAP [32] in the supplementary material.

Ablation Studies on KITTI. To study the effect of the proposed components in FUMET, we conduct ablation studies on the training losses, balancing the loss weights, camera height optimization, and outlier filtering. The results are shown in Tab. 4. Thanks to the accuracy of LSP and the robustness of Silhouette Projector, the camera height loss makes a greater contribution to the depth accuracy than the auxiliary rough geometric loss. By introducing the camera height optimization, the model can learn the scale with more accurate and consistent supervision than leveraging the size prior independently for each frame (rows 1, 2, and 6 in Tab. 4), which leads to improved accuracy. Though balancing the loss weights (Eq. (14)) with either \mathcal{L}_{aux} or \mathcal{L}_{cam} degrades the results due to the insufficient strength of the supervisory signals for metric scale, it contributes to improving accuracy when both \mathcal{L}_{aux} and \mathcal{L}_{cam} are employed.

Our proposed method can also be used with a pre-trained MDE model in an offline manner. We first train the original Monodepth2 [15] without FUMET, then predict depth maps of training images. With Silhouette Projector and LSP, we obtain per-frame *scaled* camera height and optimize it across frames with the median operation. We fix this value and use it as the camera height supervision

Table 4: Each loss contributes to learning metric scale and the others enhance the depth accuracy. Bal: balancing the loss weights (Eq. (14)). Opt: camera height optimization. Flt: outlier filtering.

\mathcal{L}_{aux}	\mathcal{L}_{cam}	Bal	Opt	Flt	AbsRel	SqRel	RMSE	$ \delta < 1.25$
✓	-	-	-	✓	0.125	0.887	5.055	0.834
✓	-	✓	-	✓	0.149	0.950	5.219	0.785
-	✓	-	✓	✓	0.116	0.875	4.874	<u>0.863</u>
-	✓	✓	-	✓	0.211	1.303	6.143	0.585
✓	✓	-	✓	✓	0.122	0.880	5.017	0.849
✓	✓	✓	-	✓	<u>0.115</u>	<u>0.822</u>	<u>4.811</u>	0.858
✓	✓	✓	✓	-	0.116	0.823	4.901	0.860
✓	✓	✓	✓	✓	0.108	0.785	4.736	0.871

Table 5: Training with a fine-tuned pseudo camera height achieves the highest accuracy. Online: optimizing the camera height while training. Offline: training with a fixed one computed from depth maps predicted with pre-trained Monodepth2. Fine-tune: training with the fixed one computed offline for 20 epochs, then optimizing online.

Deriv. method	AbsRel	SqRel	RMSE	$ \delta < 1.25$
Online	0.108	0.808	4.740	0.870
Offline	<u>0.106</u>	<u>0.790</u>	<u>4.719</u>	<u>0.873</u>
Fine-tune	0.105	0.782	4.682	0.876

for an untrained Monodepth2 on FUMET. We examine the difference between the proposed camera height optimized online and the fixed one derived offline when we use them as the pseudo-supervision. Note that we also fix λ_{cam} because the pseudo-label can be considered sufficiently accurate from the early stage of training. We also run another experiment in which we fix camera height supervision derived offline for the first 20 epochs, which is the same number of training epochs as the original Monodepth2, then unfix the camera height pseudo-label to fine-tune it via online optimization. As shown in Tab. 5, the model supervised with the fixed camera height computed offline achieves slightly better results than the one supervised with the camera height optimized online. This result indicates that our strategy of deriving *scaled* camera height is highly accurate and fixing it stabilizes the training. The model trained with the fine-tuned camera height achieves the highest accuracy, which shows that our FUMET can attain better accuracy with additional pre-training costs.

6.2 LSP

We provide implementation details of LSP in the supplementary material.

Evaluation on KITTI and Cityscapes. We evaluate the generalizability of our LSP with the dimensions of the annotated 3D bounding box on both the KITTI [12] and the Cityscapes [6] datasets in the *Car* class. We compare the results of LSP to a fixed car height prior (1.59m, following [20]) and MonoCInIS [19], a 3D object detection model that outputs object dimensions from an image and an instance mask. We train MonoCInIS from scratch with the KITTI dataset. As an evaluation metric, we compute the absolute relative error for each dimension. When we estimate car dimensions with LSP and MonoCInIS, we use a ground-truth instance mask for making a car image mask and as input, respectively. For all the experiments in LSP, input vehicle masks are cropped and resized to 300×300 resolution. On the KITTI dataset, we extract 3,769 samples for evaluation [4]. On the Cityscapes dataset, we evaluate on the validation split,

Table 6: Generalizability on KITTI and Cityscapes. We evaluate LSP on Absolute Relative Error for each dimension (height/width/length), comparing with the fixed car height prior (1.59 m, following [20]) and MonoCInIS [19] trained on KITTI. LSP robustly surpasses other methods in most cases and can predict accurately regardless of the datasets.

	Method	Easy	Moderate	Hard	All
KITTI	Fixed Height [20]	0.065 / - / -	<u>0.064</u> / - / -	0.065 / - / -	<u>0.065</u> / - / -
	MonoCInIS [19]	0.069 / 0.065 / 0.093	0.112 / 0.106 / 0.133	0.268 / 0.263 / 0.269	0.159 / 0.155 / 0.177
	LSP (only height)	<u>0.057</u> / - / -	0.089 / - / -	0.106 / - / -	0.087 / - / -
	LSP	0.046 / 0.136 / 0.130	0.060 / 0.144 / 0.139	<u>0.073</u> / 0.155 / 0.150	0.062 / 0.147 / 0.145
Cityscapes	Fixed Height [20]	0.093 / - / -	0.093 / - / -	<u>0.076</u> / - / -	<u>0.092</u> / - / -
	MonoCInIS [19]	0.578 / 0.599 / 0.584	0.759 / 0.766 / 0.767	0.843 / 0.852 / 0.852	0.851 / 0.859 / 0.856
	LSP (only height)	<u>0.063</u> / - / -	<u>0.075</u> / - / -	0.103 / - / -	0.096 / - / -
	LSP	0.056 / 0.069 / 0.067	0.068 / 0.068 / 0.078	0.069 / 0.059 / 0.075	0.080 / 0.064 / 0.077

which has 500 samples. On the KITTI dataset, each object is separated into 3 levels based on detection difficulty: Easy, Moderate, and Hard. Considering the difference in focal lengths between KITTI and Cityscapes, we also separate objects on Cityscapes into the same levels as on KITTI.

As shown in Tab. 6, LSP outperforms the other two methods in most cases. Particularly, the accuracy of MonoCInIS drops on Cityscapes since it overfits to the training dataset, *i.e.*, KITTI and cannot address the domain gap. In contrast, thanks to the large training data and various augmentations, our LSP obtains the highest generalizability and robustly estimates dimensions on any dataset. Learning not only car height but also width and length also improves the prediction accuracy as this facilitates estimation of the overall car sizes.

7 Conclusion

In this work, we proposed FUMET, a novel training framework for any MDE model to learn metric scale. FUMET aggregates the scale information obtained from a learned car size prior into the camera height and uses it as pseudo-supervision formulated as optimization across frames and epochs by leveraging its invariance. We showed that FUMET enables monocular depth estimators to embody metric scale regardless of the underlying network architectures and achieves state-of-the-art absolute-scale depth estimation accuracy on both the KITTI and the Cityscapes datasets, and demonstrated that it can leverage any dataset for training. We also proposed LSP, which can be easily trained on a large-scale dataset without manual annotations and robustly estimates vehicle dimensions from their appearances. We showed that the learned size prior is superior to a fixed one and outperforms a similar 3D object detection model. FUMET can be leveraged as a drop-in training scheme for any monocular depth estimation method. We hope it will be a versatile tool for road scene understanding research.

Acknowledgements

This work was in part supported by JSPS 20H05951 and 21H04893, JST JP-MJCR20G7 and JPMJAP2305, and RIKEN GRP.

References

1. Bhat, S.F., Alhashim, I., Wonka, P.: AdaBins: Depth Estimation Using Adaptive Bins. In: CVPR. pp. 4009–4018 (2021)
2. Casser, V., Pirk, S., Mahjourian, R., Angelova, A.: Unsupervised Monocular Depth and Ego-Motion Learning with Structure and Semantics. In: CVPRW. pp. 381–388 (2019)
3. Chawla, H., Varma, A., Arani, E., Zonooz, B.: Multimodal Scale Consistency and Awareness for Monocular Self-Supervised Depth Estimation. In: ICRA. pp. 5140–5146 (2021)
4. Chen, X., Kundu, K., Zhu, Y., Berneshawi, A.G., Ma, H., Fidler, S., Urtasun, R.: 3D Object Proposals for Accurate Object Class Detection. *Adv. Neural Inform. Process. Syst.* **28** (2015)
5. Chen, X., Li, T.H., Zhang, R., Li, G.: Frequency-Aware Self-Supervised Monocular Depth Estimation. In: WACV. pp. 5808–5817 (2023)
6. 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: CVPR (2016)
7. Ding, M., Huo, Y., Yi, H., Wang, Z., Shi, J., Lu, Z., Luo, P.: Learning Depth-Guided Convolutions for Monocular 3D Object Detection. In: CVPR. pp. 11669–11678 (2020)
8. Eigen, D., Fergus, R.: Predicting Depth, Surface Normals and Semantic Labels with a Common Multi-Scale Convolutional Architecture. In: ICCV. pp. 3213–3223 (2015)
9. Eigen, D., Puhrsch, C., Fergus, R.: Depth Map Prediction from a Single Image Using a Multi-Scale Deep Network. In: *Adv. Neural Inform. Process. Syst.* vol. 2, pp. 2366–2374 (2014)
10. Frost, D., Prisacariu, V., Murray, D.: Recovering Stable Scale in Monocular SLAM Using Object-Supplemented Bundle Adjustment. *IEEE Transactions on Robotics* **34**(3), 736–747 (2018). <https://doi.org/10.1109/TR0.2018.2820722>
11. Garg, R., Bg, V.K., Carneiro, G., Reid, I.: Unsupervised CNN for Single View Depth Estimation: Geometry to the Rescue. In: ECCV. pp. 740–756 (2016)
12. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision Meets Robotics: The KITTI Dataset. *International Journal of Robotics Research (IJRR)* **32**(11), 1231–1237 (2013). <https://doi.org/10.1177/0278364913491297>
13. Geyer, J., Kassahun, Y., Mahmudi, M., Ricou, X., Durgesh, R., Chung, A.S., Hauswald, L., Pham, V.H., Mühlegg, M., Dorn, S., Fernandez, T., Jänicke, M., Mirashi, S., Savani, C., Sturm, M., Vorobiov, O., Oelker, M., Garreis, S., Schuberth, P.: A2D2: Audi Autonomous Driving Dataset. *CoRR* **abs/2004.06320** (2020)
14. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised Monocular Depth Estimation with Left-Right Consistency. In: CVPR. pp. 6602–6611 (2017)
15. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging Into Self-Supervised Monocular Depth Estimation. In: ICCV. pp. 3827–3837 (2019)

16. Guizilini, V., Ambrus, R., Pillai, S., Raventos, A., Gaidon, A.: 3D Packing for Self-Supervised Monocular Depth Estimation. In: CVPR. pp. 2482–2491 (2020)
17. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press, ISBN: 0521540518, second edn. (2004)
18. He, M., Hui, L., Bian, Y., Ren, J., Xie, J., Yang, J.: RA-Depth: Resolution Adaptive Self-supervised Monocular Depth Estimation. In: ECCV. pp. 565–581 (2022)
19. Heylen, J., De Wolf, M., Dawagne, B., Proesmans, M., Van Gool, L., Abbeloos, W., Abdelkawy, H., Reino, D.O.: MonoCInIS: Camera Independent Monocular 3D Object Detection Using Instance Segmentation. In: ICCVW. pp. 923–934 (2021)
20. Hoiem, D., Efros, A.A., Hebert, M.: Putting Objects in Perspective. IJCV **80**, 3–15 (2008). <https://doi.org/10.1007/s11263-008-0137-5>
21. Houston, J., Zuidhof, G., Bergamini, L., Ye, Y., Chen, L., Jain, A., Omari, S., Iglovikov, V., Ondruska, P.: One Thousand and One Hours: Self-Driving Motion Prediction Dataset. In: Conference on Robot Learning. pp. 409–418 (2021)
22. Jain, J., Li, J., Chiu, M.T., Hassani, A., Orlov, N., Shi, H.: OneFormer: One Transformer to Rule Universal Image Segmentation. In: CVPR. pp. 2989–2998 (2023)
23. Kim, D., Ga, W., Ahn, P., Joo, D., Chun, S., Kim, J.: Global-Local Path Networks for Monocular Depth Estimation with Vertical CutDepth. CoRR **abs/2201.07436** (2022)
24. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. CoRR **abs/1412.6980** (2014)
25. Lyu, X., Liu, L., Wang, M., Kong, X., Liu, L., Liu, Y., Chen, X., Yuan, Y.: HR-Depth: High Resolution Self-Supervised Monocular Depth Estimation. In: AAAI. vol. 35, pp. 2294–2301 (2021)
26. Milli, E., Erkent, O., Yilmaz, A.E.: Multi-Modal Multi-Task (3MT) Road Segmentation. IEEE Robotics and Automation Letters **8**(9), 5408–5415 (2023). <https://doi.org/10.1109/LRA.2023.3295254>
27. Ning, C., Gan, H.: Trap Attention: Monocular Depth Estimation with Manual Traps. In: CVPR. pp. 5033–5043 (2023)
28. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library. In: Adv. Neural Inform. Process. Syst. pp. 8026—8037 (2019)
29. Piccinelli, L., Sakaridis, C., Yu, F.: iDisc: Internal Discretization for Monocular Depth Estimation. In: CVPR. pp. 21477–21487 (2023)
30. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet: A Large Scale Visual Recognition Challenge. IJCV **115**, 211–252 (2015)
31. Sadat, A., Casas, S., Ren, M., Wu, X., Dhawan, P., Urtasun, R.: Perceive, Predict, and Plan: Safe Motion Planning Through Interpretable Semantic Representations. In: ECCV. pp. 414–430 (2020)
32. Schonberger, J.L., Frahm, J.M.: Structure-From-Motion Revisited. In: CVPR. pp. 4104–4113 (2016)
33. Simonelli, A., Bulò, S.R., Porzi, L., Kotschieder, P., Ricci, E.: Are We Missing Confidence in Pseudo-LiDAR Methods for Monocular 3D Object Detection? In: ICCV. pp. 3225–3233 (2021)
34. Spencer, J., Russell, C., Hadfield, S., Bowden, R.: Kick Back & Relax: Learning to Reconstruct the World by Watching SlowTV (2023)

35. Sucar, E., Hayet, J.B.: Bayesian Scale Estimation for Monocular SLAM Based on Generic Object Detection for Correcting Scale Drift. In: ICRA. pp. 5152–5158 (2018)
36. Wagstaff, B., Kelly, J.: Self-Supervised Scale Recovery for Monocular Depth and Egomotion Estimation. In: IROS. pp. 2620–2627 (2021)
37. Wang, C., Buenaposada, J.M., Zhu, R., Lucey, S.: Learning Depth from Monocular Videos Using Direct Methods. In: CVPR. pp. 2022–2030 (2018)
38. Wang, H., Cai, P., Fan, R., Sun, Y., Liu, M.: End-to-End Interactive Prediction and Planning with Optical Flow Distillation for Autonomous Driving. In: CVPRW. pp. 2229–2238 (2021)
39. Wang, Y., Chao, W.L., Garg, D., Hariharan, B., Campbell, M., Weinberger, K.Q.: Pseudo-LiDAR From Visual Depth Estimation: Bridging the Gap in 3D Object Detection for Autonomous Driving. In: CVPR. pp. 8437–8445 (2019)
40. Wang, Y., Chen, X., You, Y., Li, L.E., Hariharan, B., Campbell, M., Weinberger, K.Q., Chao, W.L.: Train in Germany, Test in the USA: Making 3D Object Detectors Generalize. In: CVPR. pp. 11710–11720 (2020)
41. Wang, Y., Yang, H., Cai, J., Wang, G., Wang, J., Huang, Y.: Unsupervised Learning of Depth and Pose Based on Monocular Camera and Inertial Measurement Unit (IMU). In: ICRA. pp. 10010–10017 (2023)
42. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004). <https://doi.org/10.1109/TIP.2003.819861>
43. Watson, J., Mac Aodha, O., Prisacariu, V., Brostow, G., Firman, M.: The Temporal Opportunist: Self-Supervised Multi-Frame Monocular Depth. In: CVPR. pp. 1164–1174 (2021)
44. Wilson, B., Qi, W., Agarwal, T., Lambert, J., Singh, J., Khandelwal, S., Pan, B., Kumar, R., Hartnett, A., Pontes, J.K., Ramanan, D., Carr, P., Hays, J.: Argoverse 2: Next Generation Datasets for Self-driving Perception and Forecasting. In: NeurIPS Datasets and Benchmarks (2021)
45. Xiang, J., Wang, Y., An, L., Liu, H., Wang, Z., Liu, J.: Visual Attention-Based Self-Supervised Absolute Depth Estimation Using Geometric Priors in Autonomous Driving. *IEEE Robotics and Automation Letters* **7**(4), 11998–12005 (2022). <https://doi.org/10.1109/LRA.2022.3210298>
46. Yang, Z., Wang, P., Wang, Y., Xu, W., Nevatia, R.: LEGO: Learning Edge with Geometry All at Once by Watching Videos. In: CVPR. pp. 225–234 (2018)
47. Yin, W., Zhang, C., Chen, H., Cai, Z., Yu, G., Wang, K., Chen, X., Shen, C.: Metric3D: Towards Zero-Shot Metric 3D Prediction from a Single Image. In: ICCV. pp. 9043–9053 (2023)
48. Yin, Z., Shi, J.: GeoNet: Unsupervised Learning of Dense Depth, Optical Flow and Camera Pose. In: CVPR. pp. 1983–1992 (2018)
49. You, Y., Wang, Y., Chao, W.L., Garg, D., Pleiss, G., Hariharan, B., Campbell, M., Weinberger, K.Q.: Pseudo-LiDAR++: Accurate Depth for 3D Object Detection in Autonomous Driving. In: ICLR (2020)
50. Yuan, W., Gu, X., Dai, Z., Zhu, S., Tan, P.: Neural Window Fully-Connected CRFs for Monocular Depth Estimation. In: CVPR. pp. 3916–3925 (2022)
51. Zeng, W., Luo, W., Suo, S., Sadat, A., Yang, B., Casas, S., Urtasun, R.: End-To-End Interpretable Neural Motion Planner. In: CVPR. pp. 8652–8661 (2019)
52. Zhang, N., Nex, F., Vosselman, G., Kerle, N.: Lite-Mono: A Lightweight CNN and Transformer Architecture for Self-Supervised Monocular Depth Estimation. In: CVPR. pp. 18537–18546 (2023)

53. Zhang, S., Zhang, J., Tao, D.: Towards Scale-Aware, Robust, and Generalizable Unsupervised Monocular Depth Estimation by Integrating IMU Motion Dynamics. In: ECCV. pp. 143–160 (2022)
54. Zhang, S., Li, X., Liu, Y., Fu, H.: Scale-aware Insertion of Virtual Objects in Monocular Videos. In: ISMAR. pp. 36–44 (2020)
55. Zhou, H., Greenwood, D., Taylor, S.: Self-Supervised Monocular Depth Estimation with Internal Feature Fusion. In: BMVC (2021)
56. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised Learning of Depth and Ego-Motion from Video. In: CVPR. pp. 6612–6619 (2017)
57. Zhu, R., Yang, X., Hold-Geoffroy, Y., Perazzi, F., Eisenmann, J., Sunkavalli, K., Chandraker, M.: Single View Metrology in the Wild. In: ECCV. pp. 316–333 (2020)