

3D-PL: Domain Adaptive Depth Estimation with 3D-aware Pseudo-Labeling

Yu-Ting Yen^{1,2}, Chia-Ni Lu¹, Wei-Chen Chiu¹, Yi-Hsuan Tsai²

¹National Chiao Tung University, Taiwan ²Phiar Technologies

Abstract. For monocular depth estimation, acquiring ground truths for real data is not easy, and thus domain adaptation methods are commonly adopted using the supervised synthetic data. However, this may still incur a large domain gap due to the lack of supervision from the real data. In this paper, we develop a domain adaptation framework via generating reliable pseudo ground truths of depth from real data to provide direct supervisions. Specifically, we propose two mechanisms for pseudo-labeling: 1) 2D-based pseudo-labels via measuring the consistency of depth predictions when images are with the same content but different styles; 2) 3D-aware pseudo-labels via a point cloud completion network that learns to complete the depth values in the 3D space, thus providing more structural information in a scene to refine and generate more reliable pseudo-labels. In experiments, we show that our pseudo-labeling methods improve depth estimation in various settings, including the usage of stereo pairs during training. Furthermore, the proposed method performs favorably against several state-of-the-art unsupervised domain adaptation approaches in real-world datasets.

Keywords: domain adaptation, monocular depth estimation, pseudo-labeling

1 Introduction

Monocular depth estimation is an ill-posed problem that aims to estimate depth from a single image. Numerous supervised deep learning methods [3, 9, 11, 21, 28, 47] have made great progress in recent years. However, they need a large amount of data with ground truth depth, while acquiring such depth labels is highly expensive and time-consuming because it requires depth sensors such as LiDAR [14] or Kinect [50]. Therefore, several unsupervised methods [13, 15, 16, 30, 42, 49] have been proposed, where these approaches estimate disparity from videos or binocular stereo images without any ground truth depth. Unfortunately, since there is no strong supervision provided, unsupervised methods may not do well under situations such as occlusion or blurring in object motion. To solve this problem, recent works use synthetic datasets since the synthetic image-depth pairs are easier to obtain and have more accurate dense depth information than real-world depth maps. However, there still exists domain shift between synthetic and real datasets, and thus many works use domain adaptation [6, 23, 29, 51, 53] to overcome this issue.

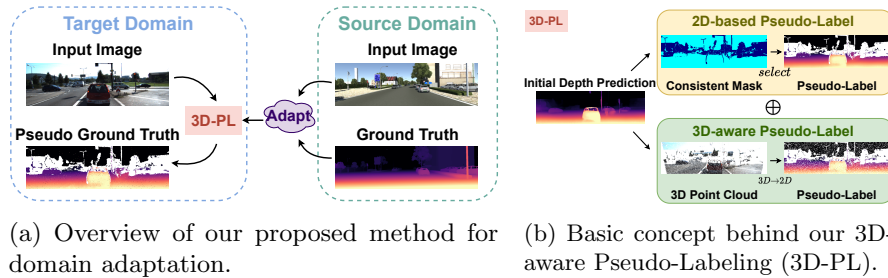


Fig. 1: **(a)** We propose a 3D-aware pseudo-labeling (3D-PL) technique to facilitate source-to-target domain adaptation for monocular depth estimation via pseudo-labeling on the target domain. **(b)** Our 3D-PL technique consists of 2D-based and 3D-aware pseudo-labels, where the former selects the pixels with highly-confident depth prediction (colorized by light blue in the consistent mask), while the latter performs 3D point cloud completion that provides refined pseudo-labels projected from 3D.

In the scenario of domain adaptation, two major techniques are commonly adopted to reduce the domain gap for depth estimation: 1) using adversarial loss [6, 23, 53] for feature-level distribution alignment, or 2) leveraging style transfer between synthetic/real data to generate real-like images as pixel-level adaptation [51, 53]. On the other hand, self-learning via pseudo-labeling the target real data is another powerful technique for domain adaptation [27, 43, 55], yet less explored in the depth estimation task. One reason is that, unlike tasks such as semantic segmentation that has the probabilistic output for classification to produce pseudo-labels, depth estimation is a regression task which requires specific designs for pseudo-label generation. In this paper, we propose novel pseudo-labeling methods in depth estimation for domain adaptation (see Fig. 1a).

To this end, we propose two mechanisms, 2D-based and 3D-aware methods, for generating pseudo depth labels (see Fig. 1b). For the 2D type, we consider the consistency of depth predictions when the model sees two images with the same content but different styles, i.e., the depth prediction can be more reliable for pixels with higher consistency. Specifically, we apply style transfer [19] to the target real image and generate its synthetic-stylized version, and then find their highly-consistent areas in depth predictions as pseudo-labels. However, this design may not be sufficient as it produces pseudo-labels only in certain confident pixels but ignore many other areas. Also, it does not take the fact that depth prediction is a 3D task into account.

To leverage the confident information obtained in our 2D-based pseudo-labeling process, we further propose to find the neighboring relationships in the 3D space via point cloud completion, so that our model is able to even select the pseudo-labels in areas that are not that confident, thus being complementary to 2D-based pseudo-labels. Specifically, we first project 2D pseudo-labels to point clouds in the 3D space, and then utilize a 3D completion model to generate

neighboring point clouds. Due to the help of more confident and accurate 2D pseudo-labels, it also facilitates 3D completion to synthesize better point clouds. Next, we project the completed point clouds back to depth values in the 2D image plane as our 3D-aware pseudo-labels. Since the 3D completion model learns the whole structural information in 3D space, it can produce reliable depth values that correct the original 2D pseudo-labels or expand extra pseudo-labels outside of the 2D ones. We also note that, although pseudo-labeling for depth has been considered in the prior work [29], different from this work that needs a pre-trained panoptic segmentation model and can only generate pseudo-labels for object instances, our method does not have this limitation as we use the point cloud completion model trained on the source domain to infer reliable 3D-aware pseudo-labels on the target image.

We conduct extensive experiments by using the virtual KITTI dataset [12] as the source domain and the KITTI dataset [14] as the real target domain. We show that both of our 2D-based and 3D-aware pseudo-labeling strategies are complementary to each other and improve the depth estimation performance. In addition, following the stereo setting in GASDA [51] where the stereo pairs are provided during training, our method can further improve the baselines and perform favorably against state-of-the-art approaches. Moreover, we directly evaluate our model on other unseen datasets, Make3D [39], and show good generalization ability against existing methods. Here are our main contributions:

- We propose a framework for domain adaptive monocular depth estimation via pseudo-labeling, consisting of 2D-based and 3D-aware strategies that are complementary to each other.
- We utilize the 2D consistency of depth predictions to obtain initial pseudo-labels, and then propose a 3D-aware method that adopts point cloud completion in the structural 3D space to refine and expand pseudo-labels.
- We show that both of our 2D-based and 3D-aware methods have advantages against existing methods on several datasets, and when having stereo pairs during training, the performance can be further improved.

2 Related Work

Monocular Depth Estimation. Monocular depth estimation is to understand 3D depth information from a single 2D image. With the recent renaissance of deep learning techniques, supervised learning methods [3, 9, 11, 21, 28, 47] have been proposed. Eigen *et al.* [9] first use a two-scale CNN-based network to directly regress on the depth, while Liu *et al.* [28] utilize continuous CRF to improve depth estimation. Furthermore, some methods propose different designs to extend the CNN-based network, such as changing the regression loss to classification [3, 11], adding geometric constraints [47], and predicting with semantic segmentation [21, 44].

Despite having promising results, the cost of collecting image-depth pairs for supervised learning is expensive. Thus, several unsupervised [13, 15, 16, 30, 42, 49] or semi-supervised [1, 17, 20, 24] methods have been proposed to estimate

disparity from the stereo pairs or videos. Garg *et al.* [13] warp the right image to reconstruct its corresponding left one (in a stereo pair) through the depth-aware geometry constraints, and take photometric error as the reconstruction penalty. Godard *et al.* [15] predict the left and right disparity separately, and enforce the left-right consistency to enhance the quality of predicted results. There are several follow-up methods to further improve the performance through semi-supervised manner [1, 24] and video self-supervision [16, 30].

Domain Adaptation for Depth Estimation. Another way to tackle the difficulty of data collection for depth estimation is to leverage the domain adaptation techniques [6, 23, 29, 34, 51, 53], where the synthetic data can provide full supervisions as the source domain and the real-world unlabeled data is the target domain. Since depth estimation is a regression task, existing methods usually rely on style transfer/image translation for pixel-level adaptation [2], adversarial learning for feature-level adaptation [23], or their combinations [51, 53]. For instance, Atapour *et al.* [2] transform the style of testing data from real to synthetic, and use it as the input to their depth prediction model that is only trained on the synthetic data. AdaDepth [23] aligns the distribution between the source and target domain at the latent feature space and the prediction level. T²net [53] further combines these two techniques, where they adopt both the synthetic-to-real translation network and the task network with feature alignment. They show that, training on the real stylized images brings promising improvement, but aligning features is not effective in the outdoor dataset.

Other follow-up methods [6, 51] take the bidirectional translation (real-to-synthetic and synthetic-to-real) and use the depth consistency loss on the prediction between the real and real-to-synthetic images. Moreover, some methods employ additional information to give constraints on the real image. GASDA [51] utilizes stereo pairs and encourages the geometry consistency to align stereo images. With a similar setting and geometry constraint to GASDA, SharinGAN [34] maps both synthetic and real images to a shared image domain for depth estimation. Moreover, DESC [29] adopts instance segmentation to apply pseudo-labeling using instance height and semantic segmentation to encourage the prediction consistency between two domains. Compared to these prior works, our proposed method provides direct supervisions on the real data in a simple and efficient pseudo-labeling way without any extra information.

Pseudo-Labeling for Depth Estimation. In general, pseudo-labeling explores the knowledge learned from labeled data to infer pseudo ground truths for unlabeled data, which is commonly used in classification [4, 18, 26, 38, 41] and scene understanding [7, 27, 32, 33, 40, 52, 55, 56] problems. Only few depth estimation methods [29, 46] adopt the concept of pseudo-labeling. DESC [29] designs a model to predict the instance height and converts the instance height to depth values as the pseudo-label for the depth prediction of the real image. Yang *et al.* [46] generate the pseudo-label from multi-view images and design a few ways to refine pseudo-labels, including fusing point clouds from multi-views. These methods succeed in producing pseudo-labels, but they require to have the instance information [29] or multi-view images [46]. Moreover, as [46]

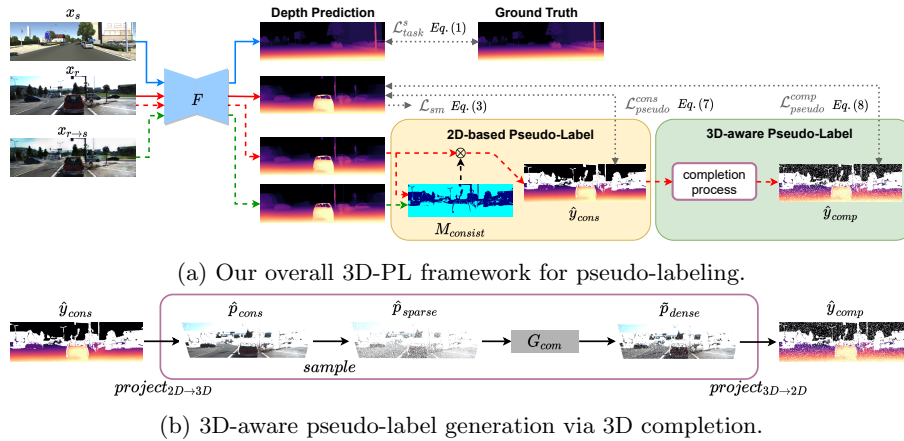


Fig. 2: **(a)** Illustration of our proposed 3D-PL framework together with the training objectives. F is the depth prediction network, with input of the synthetic image x_s , the real image x_r , and the synthetic-stylized image $x_{r \rightarrow s}$. In 3D-PL, we obtain 2D-based pseudo-labels \hat{y}_{cons} through finding the region with consistent depth (light blue color in $M_{consist}$) across the predictions of x_r and $x_{r \rightarrow s}$ (see Section 3.2), while 3D-aware pseudo-labels \hat{y}_{comp} are obtained via the 3D completion process. Here we denote solid lines as the computation flow where the gradients can be back-propagated, while the dashed lines indicate that pseudo-labels are generated offline based on the preliminary model in Section 3.1. **(b)** We first project the 2D-based pseudo-labels \hat{y}_{cons} to the 3D point cloud \hat{p}_{cons} , followed by uniformly sub-sampling \hat{p}_{cons} to sparse \hat{p}_{sparse} . Then, the completion network G_{com} densifies \hat{p}_{sparse} to obtain \tilde{p}_{dense} , in which we further project \tilde{p}_{dense} back to 2D and produce 3D-aware pseudo-labels \hat{y}_{comp} (see Section 3.2).

is a multi-stereo task, it is easier to build a complete point cloud from multi views and render the depth map as pseudo-labels. Their task also focuses on the main object instead of the overall scene. In our method, we design the point cloud completion method to generate reliable 3D-aware pseudo-labels based on a single image that contains a real-world outdoor scene.

3 Proposed Method

Our goal in this paper is to adapt the depth prediction model F to the unlabeled real image x_r as the target domain, where the synthetic image-depth pair (x_s, y_s) in the source domain is provided for supervision. Without domain adaptation, the depth prediction model F can be well trained on the synthetic data (x_s, y_s) , but it cannot directly perform well on the real image x_r because of the domain shift. Thus, we propose our pseudo-labeling method to provide direct supervisions on target image x_r , which reduces the domain gap effectively.

Fig. 2 illustrates the overall pipeline of our method. To utilize our pseudo-labeling techniques, we first use the synthetic data to train a preliminary depth prediction model F , and then adopt this pretrained model to infer pseudo-labels on real data for self-training. For pseudo-label generation, we propose 2D-based and 3D-aware schemes, where we name them as *consistency label* and *completion label*, respectively. We detail our model designs in the following sections.

3.1 Preliminary Model Objectives

Here, we describe the preliminary objectives during our model pre-training by using the synthetic image-depth pairs (x_s, y_s) and the real image x_r , including depth estimation loss and smoothness loss. Please note that, this is a common step before pseudo-labeling, in order to account for initially noisy predictions.

Depth Estimation Loss. As synthetic image-depth pairs (x_s, y_s) can provide the supervision, we directly minimize the L_1 distance between the predicted depth $\tilde{y}_s = F(x_s)$ of the synthetic image x_s and the ground truth depth y_s .

$$\mathcal{L}_{task}^s(F) = \|\tilde{y}_s - y_s\|_1. \quad (1)$$

In addition to the synthetic images x_s , we follow the similar style translation strategy as [53] to generate real-stylized images $x_{s \rightarrow r}$, in which $x_{s \rightarrow r}$ maintains the content of x_s but has the style from a randomly chosen real image x_r . Note that, to keep the simplicity of our framework, we adopt the real-time style transfer AdaIN [19] (pretrained model provided by [19]) instead of training another translation network like [53].

$$\mathcal{L}_{task}^{s \rightarrow r}(F) = \|\tilde{y}_{s \rightarrow r} - y_s\|_1. \quad (2)$$

Smoothness Loss. For the target image x_r , we adopt the smoothness loss as [15, 53] to encourage the local depth prediction \tilde{y}_r being smooth and consistent. Since depth values are often discontinuous on the boundaries of objects, we weigh this loss with the edge-aware term:

$$\mathcal{L}_{sm}(F) = e^{-\nabla x_r} \|\nabla \tilde{y}_r\|_1, \quad (3)$$

where ∇ is the first derivative along spatial directions.

3.2 Pseudo-Label Generation

With the preliminary loss functions introduced in Section 3.1 that pre-train the model, we then perform our pseudo-labeling process with two schemes. First, 2D-based consistency label aims to find the highly confident pixels from depth predictions as pseudo-labels. Second, 3D-aware completion label utilizes a 3D completion model G_{com} to refine some prior pseudo-labels and further extend the range of pseudo-labels (see Fig. 2).

2D-based Consistency Label A typical way to discover reliable pseudo-labels is to find confident ones, e.g., utilizing the softmax output from tasks like semantic segmentation [27]. However, due to the nature of the regression task in depth estimation, it is not trivial to obtain such 2D-based pseudo-labels from the network output. Therefore, we design a simple yet effective way to construct the confidence map via feeding the model two target images with the same content but different styles. Since pixels in two images have correspondence, our motivation is that, pixels that are more confident should have more consistent depth values across two predictions (i.e., finding pixels that are more domain invariant through the consistency of predictions from real images with different styles).

To achieve this, we first obtain the synthetic-stylized image $x_{r \rightarrow s}$ for the real image x_r , which combines the content of x_r and the style of a synthetic image x_s , via AdaIN [19]. Then, we obtain depth predictions of these two images, $\tilde{y}_r = F(x_r)$, $\tilde{y}_{r \rightarrow s} = F(x_{r \rightarrow s})$, and calculate their difference. If the difference at one pixel is less than a threshold τ , we consider this pixel as a more confident prediction to form the pseudo-label \hat{y}_{cons} . The procedure is written as:

$$\begin{aligned} M_{consist} &= |\tilde{y}_r - \tilde{y}_{r \rightarrow s}| < \tau, \\ \hat{y}_{cons} &= M_{consist} \otimes \tilde{y}_r, \end{aligned} \quad (4)$$

where $M_{consist}$ is the binary mask for consistency, which records where pixels are consistent. τ is the threshold, set as 0.5 in meter, and \otimes is the element-wise product to filter the prediction \tilde{y}_r of the target image.

3D-aware Completion Label Since depth estimation is a 3D problem, we expand the prior 2D-based pseudo-label \hat{y}_{cons} to obtain more pseudo-labels in the 3D space, so that the pseudo-labeling process can benefit from the learned 3D structure. To this end, based on the 2D consistency label \hat{y}_{cons} , we propose a 3D completion process to reason neighboring relationships in 3D. As shown in Fig. 2b, the 3D completion process adopts the point cloud completion technique to learn from the 3D structure and generate neighboring points.

First, we project the 2D-based pseudo-label \hat{y}_{cons} to point clouds $\hat{p}_{cons} = project_{2D \rightarrow 3D}(\hat{y}_{cons})$ in the 3D space. In the projection procedure, we reconstruct each point (x_i, y_i, z_i) from the image pixel (u_i, v_i) with its depth value d_i based on the standard pinhole camera model (more details and discussions are provided in the supplementary material). Next, we uniformly sample points from \hat{p}_{cons} to have sparse point clouds $\hat{p}_{sparse} = sample(\hat{p}_{cons})$, followed by taking \hat{p}_{sparse} as the input to the 3D completion model G_{com} for synthesizing the missing points. Those generated points from the 3D completion model G_{com} compose new dense point clouds $\tilde{p}_{dense} = G_{com}(\hat{p}_{sparse})$, and then we project each point $(\tilde{x}_i, \tilde{y}_i, \tilde{z}_i)$ back to the original 2D plane as $(\tilde{u}_i, \tilde{v}_i)$ with updated depth value $\tilde{d}_i = \tilde{z}_i$.

Therefore, our 3D-aware pseudo-label \hat{y}_{comp} (i.e., completion label) is formed by the updated depth value \tilde{d}_i . Note that, as there could exist some projected

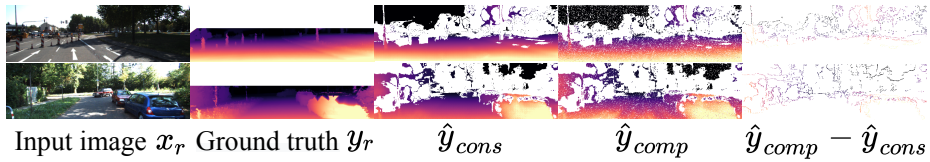


Fig. 3: Examples for our pseudo-labels. The third and fourth columns are pseudo-labels for 2D-based \hat{y}_{cons} and 3D-aware \hat{y}_{comp} . The final column represents the complementary pseudo-labels produced by \hat{y}_{comp} . Note that ground truth y_r is the reference but not used in our model training.

points falling outside the image plane and not all the pixels on the image plane are covered by the projected points, we construct a mask M_{valid} which records the pixels on the completion label \hat{y}_{comp} where the projection succeeds, i.e., having valid $(\tilde{u}_i, \tilde{v}_i)$.

$$\hat{y}_{comp} = M_{valid} \otimes project_{3D \rightarrow 2D}(\tilde{p}_{dense}^s). \quad (5)$$

In Fig. 3, we show that the 3D-aware completion label \hat{y}_{comp} expands the pseudo-labels from the 2D-based consistency label \hat{y}_{cons} , i.e., visualizations in $\hat{y}_{comp} - \hat{y}_{cons}$ are additional pseudo-labels from the 3D completion process (please refer to Section 4.3 for further analyzing the effectiveness of 3D-aware pseudo-labels).

3D Completion Model. We pre-train the 3D completion model G_{com} using the synthetic ground truth depth y_s in advance and keep it fixed during our completion process. We project the entire y_s to 3D point clouds $\hat{p}^s = project_{2D \rightarrow 3D}(y_s)$ and then perform the same process (i.e., sampling and completion) as in Fig. 2b to obtain the generated dense point clouds \tilde{p}_{dense}^s . Since \hat{p}^s is the ground truth point clouds of \tilde{p}_{dense}^s , we can directly minimize Chamfer Distance (CD) [10] between these two point clouds to train the 3D completion model G_{com} , $\mathcal{L}_{cd}(G_{com}) = CD(\hat{p}^s, \tilde{p}_{dense}^s)$.

3.3 Overall Training Pipeline and Objectives

There are two training stages in our proposed method: the first stage is to train a preliminary depth model F , and the second stage is to apply the proposed pseudo-labeling techniques through this preliminary model. The loss in the first stage consists of the ones introduced in Section 3.1:

$$\mathcal{L}_{base} = \lambda_{task}(\mathcal{L}_{task}^s + \mathcal{L}_{task}^{s \rightarrow r}) + \lambda_{sm}\mathcal{L}_{sm}, \quad (6)$$

where λ_{task} and λ_{sm} are set as 100 and 0.1 respectively, following the similar settings in [53]. Note that in our implementation, for every synthetic image x_s , we augment three corresponding real-stylized images $x_{s \rightarrow r}$, where their styles are obtained from three real images randomly drawn from the training set.

Training with Pseudo-labels. In the second stage, we use our generated 2D-based and 3D-aware pseudo-labels in Eq. (4) and Eq. (5) to provide direct supervisions on the target image x_r . Since the completion label \hat{y}_{comp} is aware of the

3D structural information and can refine the prior 2D-based pseudo-labels \hat{y}_{cons} , we choose the completion label \hat{y}_{comp} as the main reference if a pixel has both consistency label \hat{y}_{cons} and completion label \hat{y}_{comp} . The 2D and 3D pseudo-label loss functions are respectively defined as:

$$\mathcal{L}_{pseudo}^{cons}(F) = \|M'_{valid} \otimes (M_{consist} \otimes \tilde{y}_r - \hat{y}_{cons})\|_1, \quad (7)$$

$$\mathcal{L}_{pseudo}^{comp}(F) = \|M_{valid} \otimes \tilde{y}_r - \hat{y}_{comp}\|_1, \quad (8)$$

where $M'_{valid} = (1 - M_{valid})$ is the inverse mask of M_{valid} . In addition to the two pseudo-labeling objectives, we also include the supervised synthetic data to maintain the training stability. The total objective of the second stage is:

$$\begin{aligned} \mathcal{L}_{total} = & \alpha(\lambda_{cons}\mathcal{L}_{pseudo}^{cons} + \lambda_{comp}\mathcal{L}_{pseudo}^{comp}) \\ & + (1 - \alpha)\lambda_{task}^s\mathcal{L}_{task}^s + \lambda_{sm}\mathcal{L}_{sm}, \end{aligned} \quad (9)$$

where α set as 0.7 is the proportion ratio between the supervised loss of the synthetic and real image. λ_{task}^s , λ_{cons} , λ_{comp} , and λ_{sm} are set as 100, 1, 0.1, and 0.1, respectively. Here we do not include the $\mathcal{L}_{task}^{s \rightarrow r}$ loss as in Eq. (6) to make the model training more focused on the real-domain data.

Stereo Setting. The training strategy mentioned above is under the condition that we can only access the monocular single image of the real data x_r . In addition, if the stereo pairs are available during training as the setting in GASDA [51], we can further include the geometry consistency loss \mathcal{L}_{tgc} in [51] to our proposed method (more details are in the supplementary material):

$$\mathcal{L}_{stereo} = \mathcal{L}_{total} + \lambda_{tgc}\mathcal{L}_{tgc}, \quad (10)$$

where \mathcal{L}_{total} is the loss in Eq. (9), and λ_{tgc} is set as 50 following [51].

4 Experimental Results

In summary, we conduct experiments for the synthetic-to-real benchmark when only single images or stereo pairs are available during training. Then we show ablation studies to demonstrate the effectiveness of the proposed pseudo-labeling methods. Moreover, we provide discussion to validate the effectiveness of our 3D-aware pseudo-labeling method. Finally, we directly evaluate our models on two real-world datasets to show the generalization ability. More results and analysis are provided in the supplementary material.

Datasets and Evaluation Metrics. We adopt Virtual KITTI (vKITTI) [12] and real KITTI [14] as our source and target datasets respectively. vKITTI contains 21,260 synthetic image-depth pairs of the urban scene under different weather conditions. Since the maximum depth ground truth values are different in vKITTI and KITTI, we clip the maximum value to 80m as [53]. For evaluating the generalization ability, we use the KITTI Stereo [31] and Make3D [39] datasets following the prior work [51]. We use the same depth evaluation metrics as [51, 53], including four types of errors and three types of accuracy metrics.

Table 1: Quantitative results on KITTI in the single-image setting, where we denote the best results in bold. For the training data, “K”, “CS”, and “S” indicate KITTI [14], CityScapes [8], and virtual-KITTI [12] datasets respectively. We highlight the rows in gray for those methods using the domain adaptation (DA) techniques.

Method	Supervised	Dataset	Cap	Error Metrics (lower, better)				Accuracy Metrics (higher, better)		
				Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Eigen <i>et al.</i> [9]	Yes	K	80m	0.203	1.548	6.307	0.282	0.702	0.890	0.958
Liu <i>et al.</i> [28]	Yes	K	80m	0.202	1.614	6.523	0.275	0.678	0.895	0.965
Zhou <i>et al.</i> [54]	No	K	80m	0.208	1.768	6.856	0.283	0.678	0.885	0.957
Zhou <i>et al.</i> [54]	No	K+CS	80m	0.198	1.836	6.565	0.275	0.718	0.901	0.960
All synthetic	No	S	80m	0.253	2.303	6.953	0.328	0.635	0.856	0.937
All real	No	K	80m	0.158	1.151	5.285	0.238	0.811	0.934	0.970
AdaDepth [23]	No	K+S(DA)	80m	0.214	1.932	7.157	0.295	0.665	0.882	0.950
T ² Net [53]	No	K+S(DA)	80m	0.182	1.611	6.216	0.265	0.749	0.898	0.959
3D-PL (Ours)	No	K+S(DA)	80m	0.169	1.371	6.037	0.256	0.759	0.904	0.961
Garg <i>et al.</i> [13]	No	K	50m	0.169	1.080	5.104	0.273	0.740	0.904	0.962
All synthetic	No	S	50m	0.244	1.771	5.354	0.313	0.647	0.866	0.943
All real	No	K	50m	0.151	0.856	4.043	0.227	0.824	0.940	0.973
AdaDepth [23]	No	K+S(DA)	50m	0.203	1.734	6.251	0.284	0.687	0.899	0.958
T ² Net [53]	No	K+S(DA)	50m	0.168	1.199	4.674	0.243	0.772	0.912	0.966
3D-PL (Ours)	No	K+S(DA)	50m	0.162	1.049	4.463	0.239	0.776	0.916	0.968




Fig. 4: Qualitative results on KITTI in the single-image setting. We show that our 3D-PL produces more accurate results for the tree and grass (upper row) and better shapes in the car (bottom row), compared to the T²Net [53] method.

Implementation Details. Our depth prediction model F adopts the same U-net [37] structure as [53]. Following [45], the 3D completion model G_{com} is modified from PCN [48] with PointNet [35]. We implement our model based on the Pytorch framework with NVIDIA Geforce GTX 2080 Ti GPU. All networks are trained with the Adam optimizer. The depth prediction model F and 3D completion model G_{com} are trained from scratch with learning rate 10^{-4} and linear decay after 10 epochs. We train F for 20 epochs in the first stage and 10 epochs in the second stage, and pre-train G_{com} for 20 epochs. The style transfer network AdaIN is pre-trained without any finetuning. Our code and models are available at <https://github.com/ccc870206/3D-PL>.

4.1 Synthetic-to-Real Benchmark

We follow [53] to use 22,600 KITTI images from 32 scenes as the real training data, and evaluate the performance on the eigen test split [9] of 697 images from other 29 scenes. Following [51], we evaluate the depth prediction results with the ground truth depth less than 80m or 50m. There are two real-data training settings in domain adaptation for monocular depth estimation: 1) only

Table 2: Quantitative results on KITTI with having stereo pairs during training.

Method	Cap	Error Metrics (lower, better)				Accuracy Metrics (higher, better)		
		Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Synthetic + Stereo	80m	0.151	1.176	5.496	0.237	0.787	0.926	0.972
T ² Net [53] + Stereo	80m	0.154	1.115	5.504	0.233	0.800	0.929	0.971
GASDA [51] (Stereo)	80m	0.149	1.003	4.995	0.227	0.824	0.941	0.973
DESC [29] + Stereo	80m	0.122	0.946	5.019	0.217	0.843	0.942	0.974
SharinGAN [34] (Stereo)	80m	0.116	0.939	5.068	0.203	0.850	0.948	0.978
3D-PL + Stereo	80m	0.113	0.903	4.902	0.201	0.859	0.952	0.979
Synthetic + Stereo	50m	0.145	0.909	4.204	0.224	0.800	0.934	0.975
T ² Net [53] + Stereo	50m	0.148	0.828	4.123	0.219	0.815	0.938	0.975
GASDA [51] (Stereo)	50m	0.143	0.756	3.846	0.217	0.836	0.946	0.976
DESC [29] + Stereo	50m	0.116	0.725	3.880	0.206	0.855	0.948	0.976
SharinGAN [34] (Stereo)	50m	0.109	0.673	3.770	0.190	0.864	0.954	0.981
3D-PL + Stereo	50m	0.106	0.641	3.643	0.189	0.872	0.958	0.982

single real images are available and we cannot access binocular or semantic information as [53]; 2) stereo pairs are available during training, so that geometry consistency can be leveraged as [51]. Our pseudo-labeling method does not have an assumption of the data requirement, and hence we conduct experiments in these two different data settings as mentioned in Section 3.3.

Single-image Setting. In this setting, we can only access monocular real images in the whole training process, as the overall objective in Eq. (9). Table 1 shows the quantitative results, where the domain adaptation methods are highlighted in gray. “All synthetic/All real” are only trained on synthetic/real image-depth pairs, which can be viewed as the lower/upper bound. Our 3D-PL method outperforms T²Net (state-of-the-art) in every metric, especially 13% and 15% improvement in the “Sq Rel” error of 50m and 80m. Fig. 4 shows the qualitative results, where we compare our 3D-PL with T²Net [53]. In the upper row, 3D-PL produces more accurate results for the tree and grass, while T²Net predicts too far and close respectively. In the lower row, our result has a better shape in the right car and more precise depth for the left two cars.

Stereo-pair Setting. If stereo pairs are available, we can utilize the geometry constraints to have self-supervised stereo supervisions as [51] using the objective in Eq. (10). Table 2 shows that our 3D-PL achieves the best performance among state-of-the-art methods. In particular, without utilizing any other clues from real-world semantic annotation, 3D-PL outperforms DESC [29] with 12% lower “Sq Rel” error in the stereo scenario. This shows that our pseudo-labeling is able to generate more reliable pseudo-labels over the single-image setting.

Fig. 5 shows qualitative results, where we compare our 3D-PL with DESC [29] + Stereo and SharinGAN [34]. 3D-PL produces better results on the overall structure (e.g., tree, wall, and car in the top row). For challenging situations such as closer objects standing alone and hiding in a complicated farther background (e.g., road sign in the middle row, tree in the bottom row), other methods tend to produce similar depth values as the background, while 3D-PL predicts better object shape and distinguish the object from the background even if it is very thin. (e.g., traffic light in the bottom row). This shows the benefits of our 3D-aware pseudo-labeling design, which reasons the 3D structural information.

Table 3: Ablation study on KITTI in the single-image setting.

Method	Error Metrics (lower, better)			
	Abs Rel	Sq Rel	RMSE	RMSE log
Only synthetic	0.244	1.771	5.354	0.313
+ \hat{y}_{cons} (all pixels)	0.166	1.125	4.557	0.244
+ \hat{y}_{cons} (confident)	0.163	1.095	4.555	0.243
+ \hat{y}_{comp} (confident)	0.164	1.054	4.473	0.239
3D-PL (\hat{y}_{comp} all pixels)	0.161	1.070	4.504	0.240
3D-PL (\hat{y}_{comp} confident)	0.162	1.049	4.463	0.239




Fig. 5: Qualitative results on KITTI with having stereo pairs during training. We show that our 3D-PL produces better results on the overall structure (e.g., tree, wall, and car in the top row), closer objects (e.g., road sign in the middle row, tree in the bottom row), and shapes (e.g., traffic light in the bottom row), compared to DESC [29] and SharinGAN [34].

4.2 Ablation Study

We demonstrate the contributions of our model designs in Table 3 using the “50m Cap” and single-image settings, where “Only synthetic” trains only on the supervised synthetic image-depth pairs.

Importance of Pseudo-labels. First, we show that either using the 2D-based or 3D-aware pseudo-labels improve the performance, i.e., “+ \hat{y}_{cons} (confident)” and “+ \hat{y}_{comp} (confident)”. Then, our final model in “3D-PL (\hat{y}_{comp} confident)” further improves depth estimation, and shows the complementary properties of using both 2D-based and 3D-aware pseudo-labels.

Importance of Consistency Mask. We show the importance of having the consistency mask in Eq. (4) as the confidence measure. For the 2D-based pseudo-label, we compare the result of using the consistency mask “+ \hat{y}_{cons} (confident)” and the one using the entire depth prediction as the pseudo-label, “+ \hat{y}_{cons} (all pixels)”. With the consistency mask, it has 3% lower in the “Sq Rel” error. Moreover, this consistency mask also improves 3D-aware pseudo-labeling when we project depth values to point clouds for 3D completion. When inputting all the pixels for this process, i.e., “3D-PL (\hat{y}_{comp} all pixels)”, this may include less accurate depth values for performing 3D completion, which results in less reliable pseudo-labels compared to our final model using the confident pixels, i.e., “3D-PL (\hat{y}_{comp} confident)”.

Table 4: Statistics of pixel proportion in our 2D/3D pseudo-labels. ‘‘R’’ and ‘‘E’’ indicate ‘‘refined’’ and ‘‘extended’’.

Method	2D Proportion	3D Proportion
2D only (+ \hat{y}_{cons})	48.91%	0%
3D-PL	5.28%	43.63% (R) + 3.9% (E)

Table 5: Results of using either the 2D-based \hat{y}_{cons} or the 3D-aware \hat{y}_{comp} pseudo-label as the reference, when there is a duplication on both pseudo-labels.

Main Reference	Error Metrics (lower, better)			
	Abs Rel	Sq Rel	RMSE	RMSE log
Completion Label \hat{y}_{comp}	0.162	1.049	4.463	0.239
Consistency Label \hat{y}_{cons}	0.164	1.095	4.529	0.243

4.3 Effectiveness of 3D-aware Pseudo-labels

To show the impact of 3D-aware pseudo-labels, we compute the proportions of pixels chosen as 2D/3D pseudo-labels in each image and take the average as the final statistics. The effectiveness of 3D-aware pseudo-labels is in two-fold: **refine** and **extend** from 2D-based pseudo-labels. In Table 4, the initial proportion of confident 2D-based pseudo-labels ‘‘2D only (+ \hat{y}_{cons})’’ is 48.91% among image pixels. As stated in Section 3.3, 3D-PL improves original 2D labels, which results in 43.63% refined and 3.9% extended labels. The rightmost subfigure of Fig. 3 visualizes extended labels $\hat{y}_{comp} - \hat{y}_{cons}$, in which it shows that the improved performance is contributed by the larger proportion of 3D-aware pseudo-labels. **Ability of pseudo-label refinement.** Since the 2D-based and 3D-aware pseudo-labels may have the duplication on the same pixel, we conduct experiments to use either \hat{y}_{cons} or \hat{y}_{comp} as the reference when such cases happen. In Table 5, choosing \hat{y}_{comp} as the main reference has the better performance, which indicates that updating the pseudo-label of a pixel from original \hat{y}_{cons} to \hat{y}_{comp} brings the positive effect. This validates that \hat{y}_{comp} can refine the prior 2D-based pseudo-labels since it is aware of the 3D structural information.

4.4 Generalization to Real-world Datasets

KITTI Stereo. We evaluate our model on 200 images of KITTI stereo 2015 [31], which is a small subset of KITTI images but has different ways of collecting groundtruth of depth information. Since the ground truth of KITTI stereo has been optimized for the moving objects, it is denser than LiDAR, especially for the vehicles. Note that, this benefits DESC [29] in this evaluation as their method relies on the instance information from the pre-trained segmentation model. Table 6 shows the quantitative results, where our 3D-PL in both single-image and stereo settings performs competitively against existing methods.

Make3D Dataset. Moreover, we directly evaluate the model on the Make3D dataset [39] without any finetuning. We choose 134 test images with central image crop and clamp the depth value to 70m, following [15]. Here, since Make3D

Table 6: Quantitative results on KITTI stereo 2015 benchmark [31]. “S^o” denotes synthetic data that [2] captures from GTA [36]. “Supervised” represents whether the method is trained on KITTI stereo.

Method	Supervised	Dataset	Error Metrics (lower, better)				Accuracy Metrics (higher, better)		
			Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Godard <i>et al.</i> [15]	No	K	0.124	1.388	6.125	0.217	0.841	0.936	0.975
Godard <i>et al.</i> [15]	No	K+CS	0.104	1.070	5.417	0.188	0.875	0.956	0.983
Atapour <i>et al.</i> [2]	No	K+S ^o (DA)	0.101	1.048	5.308	0.184	0.903	0.988	0.992
T ² Net [53]	No	K+S(DA)	0.155	1.731	6.510	0.237	0.800	0.921	0.969
3D-PL	No	K+S(DA)	0.147	1.352	6.157	0.233	0.800	0.918	0.967
GASDA [51] (Stereo)	No	K+S(DA)	0.106	0.987	5.215	0.176	0.885	0.963	0.986
DESC [29] + Stereo	No	K+S(DA)	0.085	0.781	4.490	0.158	0.909	0.967	0.986
SharinGAN [34] (Stereo)	No	K+S(DA)	0.092	0.904	4.614	0.159	0.906	0.969	0.987
3D-PL + Stereo	No	K+S(DA)	0.085	0.830	4.489	0.149	0.915	0.971	0.988

Table 7: Quantitative results on Make3D [39]. “Supervised” represents whether the method is trained on Make3D.

Method	Supervised	Error Metrics (lower, better)		
		Abs Rel	Sq Rel	RMSE
Karsch <i>et al.</i> [22]	Yes	0.398	4.723	7.801
Laina <i>et al.</i> [25]	Yes	0.198	1.665	5.461
AdaDepth [23]	Yes	0.452	5.71	9.559
Godard <i>et al.</i> [15]	No	0.505	10.172	10.936
AdaDepth [23]	No	0.647	12.341	11.567
T ² Net [53]	No	0.508	6.589	8.935
Atapour <i>et al.</i> [2]	No	0.423	9.343	9.002
GASDA [51]	No	0.403	6.709	10.424
DESC [29]	No	0.393	4.604	8.126
SharinGAN [34]	No	0.377	4.900	8.388
S2R-DepthNet [5]	No	0.490	10.681	10.892
3D-PL	No	0.352	3.539	7.967

is a different domain from the KITTI training data, we apply the single-image model to reduce the strong domain-related constraints such as the stereo supervisions. In Table 7, 3D-PL achieves the best performance compared to other approaches. It is also worth mentioning that 3D-PL outperforms the domain generalization method [5] and supervised method [22] by 66% and 25% in “Sq Rel”, showing the promising generalization capability.

5 Conclusions

In this paper, we introduce a domain adaptation method for monocular depth estimation. We propose 2D-based and 3D-aware pseudo-labeling mechanisms, which utilize knowledge from synthetic domain as well as 3D structural information to generate reliable pseudo depth labels for real data. Extensive experiments show that our pseudo-labeling strategies are able to improve depth estimation in various settings against several state-of-the-art domain adaptation approaches, as well as achieving good performance in unseen datasets for generalization. **Acknowledgement.** This project is supported by MOST (Ministry of Science and Technology, Taiwan) 111-2636-E-A49-003 and 111-2628-E-A49-018-MY4.

References

1. Amiri, A.J., Loo, S.Y., Zhang, H.: Semi-supervised monocular depth estimation with left-right consistency using deep neural network. In: IEEE International Conference on Robotics and Biomimetics (ROBIO) (2019) [3](#), [4](#)
2. Atapour-Abarghouei, A., Breckon, T.P.: Real-time monocular depth estimation using synthetic data with domain adaptation via image style transfer. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [4](#), [14](#)
3. Cao, Y., Wu, Z., Shen, C.: Estimating depth from monocular images as classification using deep fully convolutional residual networks. IEEE Transactions on Circuits and Systems for Video Technology (TCSVT) (2017) [1](#), [3](#)
4. Chen, C., Xie, W., Huang, W., Rong, Y., Ding, X., Huang, Y., Xu, T., Huang, J.: Progressive feature alignment for unsupervised domain adaptation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [4](#)
5. Chen, X., Wang, Y., Chen, X., Zeng, W.: S2r-depthnet: Learning a generalizable depth-specific structural representation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [14](#)
6. Chen, Y.C., Lin, Y.Y., Yang, M.H., Huang, J.B.: Crdoco: Pixel-level domain transfer with cross-domain consistency. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [1](#), [2](#), [4](#)
7. Chen, Z., Zhang, R., Zhang, G., Ma, Z., Lei, T.: Digging into pseudo label: a low-budget approach for semi-supervised semantic segmentation. IEEE Access (2020) [4](#)
8. 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: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) [10](#)
9. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. ArXiv:1406.2283 (2014) [1](#), [3](#), [10](#)
10. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [8](#)
11. Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [1](#), [3](#)
12. Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) [3](#), [9](#), [10](#)
13. Garg, R., Bg, V.K., Carneiro, G., Reid, I.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: European Conference on Computer Vision (ECCV) (2016) [1](#), [3](#), [4](#), [10](#)
14. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2012) [1](#), [3](#), [9](#), [10](#)
15. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [1](#), [3](#), [4](#), [6](#), [13](#), [14](#)
16. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: IEEE International Conference on Computer Vision (ICCV) (2019) [1](#), [3](#), [4](#)

17. Guizilini, V., Li, J., Ambrus, R., Pillai, S., Gaidon, A.: Robust semi-supervised monocular depth estimation with reprojected distances. In: Conference on Robot Learning (CoRL) (2020) [3](#)
18. Hu, Z., Yang, Z., Hu, X., Nevatia, R.: Simple: Similar pseudo label exploitation for semi-supervised classification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
19. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: IEEE International Conference on Computer Vision (ICCV) (2017) [2](#), [6](#), [7](#)
20. Ji, R., Li, K., Wang, Y., Sun, X., Guo, F., Guo, X., Wu, Y., Huang, F., Luo, J.: Semi-supervised adversarial monocular depth estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2019) [3](#)
21. Jiao, J., Cao, Y., Song, Y., Lau, R.: Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss. In: European Conference on Computer Vision (ECCV) (2018) [1](#), [3](#)
22. Karsch, K., Liu, C., Kang, S.B.: Depth transfer: Depth extraction from video using non-parametric sampling. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2014) [14](#)
23. Kundu, J.N., Uppala, P.K., Pahuja, A., Babu, R.V.: Adadepth: Unsupervised content congruent adaptation for depth estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [1](#), [2](#), [4](#), [10](#), [14](#)
24. Kuznetsov, Y., Stuckler, J., Leibe, B.: Semi-supervised deep learning for monocular depth map prediction. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [3](#), [4](#)
25. Laina, I., Rupprecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks. In: International Conference on 3D Vision (3DV) (2016) [14](#)
26. Lee, D.H., et al.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: International Conference on Machine Learning (ICML) (2013) [4](#)
27. Li, Y., Yuan, L., Vasconcelos, N.: Bidirectional learning for domain adaptation of semantic segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [2](#), [4](#), [7](#)
28. Liu, F., Shen, C., Lin, G., Reid, I.: Learning depth from single monocular images using deep convolutional neural fields. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2015) [1](#), [3](#), [10](#)
29. Lopez-Rodriguez, A., Mikolajczyk, K.: Desc: Domain adaptation for depth estimation via semantic consistency. ArXiv:2009.01579 (2020) [1](#), [3](#), [4](#), [11](#), [12](#), [13](#), [14](#)
30. Mahjourian, R., Wicke, M., Angelova, A.: Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [1](#), [3](#), [4](#)
31. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015) [9](#), [13](#), [14](#)
32. Pastore, G., Cermelli, F., Xian, Y., Mancini, M., Akata, Z., Caputo, B.: A closer look at self-training for zero-label semantic segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
33. Paul, S., Tsai, Y.H., Schuster, S., Roy-Chowdhury, A.K., Chandraker, M.: Domain adaptive semantic segmentation using weak labels. In: European Conference on Computer Vision (ECCV) (2020) [4](#)

34. PNVr, K., Zhou, H., Jacobs, D.: Sharingan: Combining synthetic and real data for unsupervised geometry estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [4](#), [11](#), [12](#), [14](#)
35. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [10](#)
36. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: European Conference on Computer Vision (ECCV) (2016) [14](#)
37. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI) (2015) [10](#)
38. Saito, K., Ushiku, Y., Harada, T.: Asymmetric tri-training for unsupervised domain adaptation. In: International Conference on Machine Learning (ICML) (2017) [4](#)
39. Saxena, A., Sun, M., Ng, A.Y.: Make3d: Learning 3d scene structure from a single still image. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2008) [3](#), [9](#), [13](#), [14](#)
40. Shin, I., Tsai, Y.H., Zhuang, B., Schulter, S., Liu, B., Garg, S., Kweon, I.S., Yoon, K.J.: Mm-tta: Multi-modal test-time adaptation for 3d semantic segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022) [4](#)
41. Taherkhani, F., Dabouei, A., Soleymani, S., Dawson, J., Nasrabadi, N.M.: Self-supervised wasserstein pseudo-labeling for semi-supervised image classification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
42. Tosi, F., Aleotti, F., Poggi, M., Mattoccia, S.: Learning monocular depth estimation infusing traditional stereo knowledge. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [1](#), [3](#)
43. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [2](#)
44. Wang, P., Shen, X., Lin, Z., Cohen, S., Price, B., Yuille, A.L.: Towards unified depth and semantic prediction from a single image. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015) [3](#)
45. Xiang, R., Zheng, F., Su, H., Zhang, Z.: 3ddepthnet: Point cloud guided depth completion network for sparse depth and single color image. ArXiv:2003.09175 (2020) [10](#)
46. Yang, J., Alvarez, J.M., Liu, M.: Self-supervised learning of depth inference for multi-view stereo. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
47. Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction. In: IEEE International Conference on Computer Vision (ICCV) (2019) [1](#), [3](#)
48. Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: Pcn: Point completion network. In: International Conference on 3D Vision (3DV) (2018) [10](#)
49. Zhan, H., Garg, R., Weerasekera, C.S., Li, K., Agarwal, H., Reid, I.: Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [1](#), [3](#)
50. Zhang, Z.: Microsoft kinect sensor and its effect. IEEE multimedia (2012) [1](#)
51. Zhao, S., Fu, H., Gong, M., Tao, D.: Geometry-aware symmetric domain adaptation for monocular depth estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [1](#), [2](#), [3](#), [4](#), [9](#), [10](#), [11](#), [14](#)

52. Zhao, X., Schuster, S., Sharma, G., Tsai, Y.H., Chandraker, M., Wu, Y.: Object detection with a unified label space from multiple datasets. In: European Conference on Computer Vision (ECCV) (2020) [4](#)
53. Zheng, C., Cham, T.J., Cai, J.: T2net: Synthetic-to-realistic translation for solving single-image depth estimation tasks. In: European Conference on Computer Vision (ECCV) (2018) [1](#), [2](#), [4](#), [6](#), [8](#), [9](#), [10](#), [11](#), [14](#)
54. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [10](#)
55. Zou, Y., Yu, Z., Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: European Conference on Computer Vision (ECCV) (2018) [2](#), [4](#)
56. Zou, Y., Zhang, Z., Zhang, H., Li, C.L., Bian, X., Huang, J.B., Pfister, T.: Pseudoseg: Designing pseudo labels for semantic segmentation. ArXiv:2010.09713 (2020) [4](#)