Learning to Relate Depth and Semantics for Unsupervised Domain Adaptation

Conference Paper**Author(s):**

Saha, Suman; Obukhov, Anton; Paudel, Danda Pani; Kanakis, Menelaos; [Chen, Yuhua](#) ; Georgoulis, Stamatios; Van Gool, Luc

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000517369>

Rights / license:

[Creative Commons Attribution-ShareAlike 4.0 International](#)

Originally published in:

<https://doi.org/10.1109/CVPR46437.2021.00810>

Learning to Relate Depth and Semantics for Unsupervised Domain Adaptation

Suman Saha*
ETH Zurich

Anton Obukhov*
ETH Zurich

Danda Pani Paudel
ETH Zurich

Menelaos Kanakis
ETH Zurich

Yuhua Chen
ETH Zurich

Stamatios Georgoulis
ETH Zurich

Luc Van Gool
ETH Zurich, KU Leuven

Abstract

We present an approach for encoding visual task relationships to improve model performance in an Unsupervised Domain Adaptation (UDA) setting. Semantic segmentation and monocular depth estimation are shown to be complementary tasks; in a multi-task learning setting, a proper encoding of their relationships can further improve performance on both tasks. Motivated by this observation, we propose a novel Cross-Task Relation Layer (CTRL), which encodes task dependencies between the semantic and depth predictions. To capture the cross-task relationships, we propose a neural network architecture that contains task-specific and cross-task refinement heads. Furthermore, we propose an Iterative Self-Learning (ISL) training scheme, which exploits semantic pseudo-labels to provide extra supervision on the target domain. We experimentally observe improvements in both tasks' performance because the complementary information present in these tasks is better captured. Specifically, we show that: (1) our approach improves performance on all tasks when they are complementary and mutually dependent; (2) the CTRL helps to improve both semantic segmentation and depth estimation tasks performance in the challenging UDA setting; (3) the proposed ISL training scheme further improves the semantic segmentation performance. The implementation is available at <https://github.com/susaha/ctrl-uda>.

1. Introduction

Semantic segmentation and monocular depth estimation are two important computer vision tasks that allow us to perceive the world around us and enable agents' reasoning, e.g., in an autonomous driving scenario. Moreover, these tasks have been shown to be complementary to each other, i.e., information from one task can improve the other task's performance [27, 40, 57]. Domain Adaptation (DA) [11]

Figure 1: Semantic segmentation improvement with our approach to unsupervised domain adaptation over the state-of-the-art DADA [59] method. Left to right: Cityscapes test images, DADA, and the proposed method (CTRL). Our model correctly segments the “bus”, “rider”, and “wall” classes underrepresented in the target domain (highlighted).

refers to maximizing model performance in an environment with a smaller degree of supervision (the *target domain*) relative to what the model was trained on (the *source domain*). Unsupervised Domain Adaptation (UDA) assumes only access to the unannotated samples from the target domain at train time – the setting of interest in this paper.

Recent domain adaptation techniques [32, 59] proposed to leverage depth information available in the source domain to improve semantic segmentation on the target domain. However, they lack an explicit multi-task formulation to relate depth and semantics, that is to say, *how each semantic category relates to different depth levels*. The term *depth levels* refers to different discrete ranges of depth values, i.e., “near” (1-5m); “medium-range” (5-20m), or “far” (>20m). This paper aims to design a model that learns explicit relationships between different visual semantic classes and depth levels within the UDA context.

To this end, we design a network architecture and a new multitask-aware feature space alignment mechanism for UDA. First, we propose a Cross-Task Relation Layer

Corresponding author: Suman Saha (suman.saha@vision.ee.ethz.ch)

* Equal contribution.

(CTRL) – a novel parameter-free differentiable module tailored to capture the task relationships given the network’s semantic and depth predictions. Second, we utilize a Semantics Refinement Head (SRH) that explicitly captures cross-task relationships by learning to predict semantic segmentation given predicted depth features. Both CTRL and SRH boost the model’s ability to effectively encode correlations between semantics and depth, thus improving predictions on the target domain. Third, we employ an Iterative Self Learning (ISL) scheme. Coupled with the model design, it further pushes the performance of semantic segmentation. As a result, our method achieves state-of-the-art semantic segmentation performance on three challenging UDA benchmarks (Sec. 4). Fig. 1 demonstrates our method’s effectiveness by comparing semantic predictions of classes underrepresented in the target domain to predictions made by the previous state-of-the-art method. The paper is organized as follows: Sec. 2 discusses the related work; Sec. 3 describes the proposed approach to UDA, the network architecture, and the learning scheme; Sec. 4 presents the experimental analysis with ablation studies; Sec. 5 concludes the paper.

2. Related Work

Semantic Segmentation. refers to the task of assigning a semantic label to each pixel of an image. Conventionally, the task has been addressed using hand-crafted features combined with classifiers, such as Random Forests [51], SVMs [16], or Conditional Random Fields [29]. Powered by the effectiveness of Convolutional Neural Networks (CNNs) [31], we have seen an increasing number of deep learning-based models. Long et al. [36] were among the first to use fully convolutional networks (FCNs) for semantic segmentation. Since then, this design has quickly become a state-of-the-art method for the task. The encoder-decoder design is still widely used [64, 5, 1, 69, 4].

Cross-domain Semantic Segmentation. Training deep networks for semantic segmentation requires large amounts of labeled data, which presents a significant bottleneck in practice, as acquiring pixel-wise labels is a labor-intensive process. A common approach to address the issue is to train the model on a source domain and apply it to a target domain in a UDA context. However, this often causes a performance drop due to the domain shift. Domain Adaptation aims to solve the issue by aligning the features from different domains. DA is a highly active research field, and techniques have been developed for various applications, including image classification [17, 34, 37, 38], object detection [8], fine-grained recognition [18], etc.

More related to our method are several works on unsupervised domain adaptation for semantic segmentation [66, 50, 71, 9, 58, 23, 62, 70, 46, 63]. This problem has been tackled with curriculum learning [66], GANs [50], adver-

sarial training on the feature space [9], output space [53], or entropy maps [58], self-learning using pseudo- or weak labels [71, 46, 23]. However, prior works typically only consider adapting semantic segmentation while neglecting any multi-task correlations. A few methods [7, 59] model correlations between semantic segmentation and depth estimation, similarly to our work, yet – as explained in Sec. 1 – these works come with crucial limitations.

Monocular Depth Estimation. Similar to semantic segmentation, monocular depth estimation is dominated by CNN-based methods [13, 15, 30, 33]. [13] introduced a CNN-based architecture for depth estimation, which regresses a dense depth map. Their approach was then improved by incorporating techniques such as a CRF [35, 33] and multi-scale CRF techniques [61]. Besides, improvements in the loss design itself also lead to better depth estimation. Examples include the reverse Huber (berHu) loss [44, 72], and the ordinal regression loss [15].

Multi-task Learning for Semantic Segmentation and Depth Estimation. Within the context of multi-task learning, semantic segmentation is shown to be highly correlated with depth estimation, and vice versa [65, 60, 27, 67, 68, 40, 52, 57, 56, 26]. To leverage this correlation, some authors have proposed to learn them jointly [48, 25, 6]. In particular, [43, 25, 55, 3] proposed to share the encoder and use multiple decoders, whereas a shared conditional decoder is used in [6]. Semantic segmentation was also demonstrated to help guide the depth training process [19, 24].

In this paper, we build upon these observations. We argue that task relationships, like the ones between depth and semantics, are not entirely domain-specific. As a result, if we correctly model these relationships in one domain, they can be transferred to another domain to help guide the DA process. The proposed method and its components are explicitly designed around this hypothesis.

3. Method

In this section, we describe our approach to UDA in the autonomous driving setting. Sec. 3.1 presents an overview of the proposed approach; Sec. 3.2 explains the notation and problem formulation; Sec. 3.3 describes supervision on the source domain; Sec. 3.4 presents the CTRL module design; Sec. 3.5 describes the ISL technique; Sec. 3.6 prescribes the rest of the network architecture details.

3.1. Overview

The primary hypothesis behind our approach is that task dependencies persist across domains, i.e., most semantic classes fall under a finite depth range. We can exploit this information from source samples and transfer it to target using adversarial training. As our goal is to train the network in a UDA setting, we follow an adversarial training scheme [22, 53] to learn domain invariant representations.

Figure 2: Overview of the proposed neural architecture (Sec. 3.2) and the CTRL module (Sec. 3.4). Supervised losses (in the middle) are applied only on the source domain; the rest of the data flow is domain-agnostic. Legend: **learned modules**, **predictions**, **loss functions**; rounded corners denote operators, rectangles denote activations.

Unlike [59] that directly aligns a combination of semantics and depth features, we wish to design a joint feature space for domain alignment by fusing the task-specific and the cross-task features and then learn to minimize the domain gap through adversarial training. To this end, we propose CTRL – a novel module that constructs the joint feature space by computing entropy maps of both the semantic label and discretized depth distributions (Fig. 2). Thus, CTRL entropy maps, generated on the source and target domains, are expected to carry similar information.

Further enhancement of semantic segmentation performance appears possible by utilizing the Iterative Self-Learning (ISL) training scheme, which does not require expensive patch-based pseudo-label generation like [23]. As our CTRL helps the network to predict high-quality predictions (Fig. 1), ISL training exploits high-confidence predictions as supervision (pseudo-labels) on the target domain.

3.2. Problem Formulation

Let $\mathcal{D}^{(s)}$ and $\mathcal{D}^{(t)}$ denote the source and target domains, with samples from them represented by tuples $(x^{(s)}, y^{(s)}, z^{(s)})$ and $(x^{(t)}, y^{(t)}, z^{(t)})$ respectively, where $x \in \mathbb{R}^{H \times W \times 3}$ are color images, $y \in \{1, \dots, C\}^{H \times W}$ are semantic annotations with C classes, and $z \in [Z_{\min}, Z_{\max}]^{H \times W}$ are depth maps from a finite frustum. Furthermore, $\mathcal{F}_e$ is the shared feature extractor, which includes a pretrained backbone, and a decoder; $\mathcal{F}_s$ and $\mathcal{F}_d$ are the task-specific semantics and depth heads, respectively; $\mathcal{F}_r$ is the SRH (Fig. 2).

First, $\mathcal{F}_e$ extracts a shared feature map to be used by SRH and task-specific semantics and depth heads. The semantics head $\mathcal{F}_s$ predicts a semantic segmentation map $\hat{y}_s = \mathcal{F}_s(\mathcal{F}_e(x))$ with C channels per pixel, denoting predicted class probabilities. The depth head $\mathcal{F}_d$ predicts a real-valued depth map $\hat{z} = \mathcal{F}_d(\mathcal{F}_e(x))$, where each pixel is mapped into the finite frustum specified in the source domain. We further employ SRH to learn the cross-task relationship between semantics and depth by making it predict semantics from the shared feature map, attenuated by the predicted depth map. Formally, the shared feature map is

point-wise multiplied by the predicted depth map, and then SRH predicts a second (auxiliary) semantic segmentation map: $\hat{y}_r = \mathcal{F}_r(\hat{z} \odot \mathcal{F}_e(x))$.

We refer to the part of the model enclosing the $\mathcal{F}_e, \mathcal{F}_s, \mathcal{F}_r, \mathcal{F}_d$ modules as a prediction network. The predictions made by the network on the source and target domains are denoted as $(\hat{y}_s^{(s)}, \hat{y}_r^{(s)}, \hat{z}^{(s)})$ and $(\hat{y}_s^{(t)}, \hat{y}_r^{(t)}, \hat{z}^{(t)})$, respectively. We upscale these predictions along the spatial dimension to match the original input image dimension $H \times W$ before any further processing. Given these semantics and depth predictions on the source and target domains, we optimize the network cost using supervised loss on the source domain, and unsupervised domain alignment loss on the target domain within the same training process.

3.3. Supervised Learning

Since the semantic segmentation predictions $\hat{y}_s^{(s)}, \hat{y}_r^{(s)}$ and ground truth $y^{(s)}$ are represented as pixel-wise class probabilities over C classes, we employ the standard cross-entropy loss with the semantic heads:

$$\mathcal{L}_{CE}(\hat{y}, y) = - \sum_{i=1}^C y_i \log \hat{y}_i. \quad (1)$$

We use the berHu loss (the reversed Huber criterion [30]) for penalizing depth predictions:

$$\mathcal{L}_{berHu}(\hat{z}, z) = \begin{cases} |z - \hat{z}| & |z - \hat{z}| \leq L, \\ \frac{(z - \hat{z})^2 + L^2}{2L} & |z - \hat{z}| > L, \end{cases} \quad (2)$$

$$L = 0.2 \max(|z - \hat{z}|).$$

Following [27], we regress inverse depth values (normalized disparity), which is shown to improve the precision of predictions on the full range of the view frustum. The parameters of the network $\theta_e, \theta_s, \theta_r, \theta_d$ (parameterizing $\mathcal{F}_e, \mathcal{F}_s, \mathcal{F}_r, \mathcal{F}_d$ modules), collectively denoted as θ_{net} , are learned to minimize the following supervised objective on

the source domain:

$$\min_{\theta_{\text{net}}} \mathbb{E}_{\mathcal{D}^{(s)}} [\mathcal{L}_{\text{CE}}(\hat{y}_s^{(s)}, y^{(s)}) + \lambda_r \mathcal{L}_{\text{CE}}(\hat{y}_r^{(s)}, y^{(s)}) + \lambda_d \mathcal{L}_{\text{berHu}}(\hat{z}^{(s)}, z^{(s)})] \quad (3)$$

where λ_r and λ_d are the hyperparameters weighting relative importance of the SRH and depth supervision.

3.4. Cross-Task Relation Layer

In the absence of ground truth annotations for the target samples, we train the network on the target images using unsupervised domain alignment loss. Existing works either align source and target domain in a semantic space [58] or a depth-aware semantic space [59] by fusing the continuous depth predictions with predicted semantic maps. Here, we argue that simple fusion of the continuous depth prediction into the semantics does not enable the network to learn useful semantic features at different depth levels. Instead, explicit modeling is required to achieve this goal.

Humans learn to relate semantic categories at each discrete depth level differently. For example, “sky” is “far away” (large depth), “vehicles” are “nearby”, “road” appears to be both “far” and “nearby”. Taking inspiration from the way humans relate semantic and depth, we design a CTRL (Fig. 2) that captures the semantic class-specific dependencies at different discrete depth levels. Moreover, CTRL also preserves task-specific information by fusing task-specific and task-dependent features learned by the semantics, depth, and refinement (SRH) heads. CTRL consists of a depth discretization, an entropy map generation, and a fusion layer described in the following subsections.

3.4.1 Depth Discretization Module

The prediction made by the depth head $\hat{z}$ contains continuous depth values. We want to map it to a discrete probability space to learn visual semantic features at different depth levels. We quantize the view frustum depth range into a set of representative discrete values following the spacing-increasing discretization (SID) [15]. Such discretization assigns progressively large depth sub-ranges further away from the point of view into separate bins, which allows us to simulate the human perception of depth relations in the scene, with a finite number of categories.

Given the depth range $[Z_{\min}, Z_{\max}]$ and the number of depth bins K , SID outputs a K -dimensional vector of discretization bin centers b as follows:

$$b_i = Z_{\min}^{1-(2i+1)/2K} \cdot Z_{\max}^{(2i+1)/2K}, \quad i = 0, \dots, K-1 \quad (4)$$

We can now assign probabilities of the predicted depth values falling into the defined bins:

$$\hat{z}' = \text{softmax}(-(\hat{z} - b)^2). \quad (5)$$

3.4.2 Joint Space for Domain Alignment

The task-dependency $\hat{y}_r$ (output by SRH), alongside the task-specific semantics $\hat{y}_s$ and depth $\hat{z}'$ probability maps, can be considered as discrete distributions over semantic classes and depth levels. As we do not have access to the ground truth labels for the target domain, one way to train the network to predict high-confidence predictions is by minimizing the uncertainty (or entropy) in the predicted distributions over the target domain [58]. The source and target domains share similar spatial features, and it is recommended to align them in the structured output space [21].

To this end, we propose a novel UDA training scheme, where task-specific and task-dependent knowledge is transferred from the source to the target domain by constraining the target distributions to be similar to the source by aligning the entropy maps of $\hat{y}_r$, $\hat{y}_s$, and $\hat{z}'$. Note that unlike [59, 58], which constrain only on the task-specific space ($\hat{y}_s$ in our case) for domain alignment, we train the network to output highly certain predictions by aligning features in the task-specific and task-dependent spaces.

We argue that aligning source and target distributions jointly in task-specific and task-dependent spaces helps to bridge the domain gap for underrepresented classes, which are learned poorly without the presence of a joint representation. To encode such a joint representation, we generate entropy maps as follows:

$$\mathcal{E}(p) = -p \odot \log p \\ E_r = \mathcal{E}(\hat{y}_r), \quad E_s = \mathcal{E}(\hat{y}_s), \quad E_d = \mathcal{E}(\hat{z}'). \quad (6)$$

We then concatenate these maps along the channel dimension to get the fused entropy map $E = \text{concat}(E_r, E_s, E_d)$ and employ adversarial training on it.

For aligning the source and target domain distributions, we train the proposed segmentation and depth prediction network (parameterized by θ_{net}) and the discriminator network $\mathcal{D}$ (parameterized by $\theta_{\mathcal{D}}$) following an adversarial learning scheme. More specifically, the discriminator is trained to correctly classify the sample domain being either source or target given only the fused entropy map:

$$\min_{\theta_{\mathcal{D}}} \left\{ \mathbb{E}_{\mathcal{D}^{(s)}} [\log \mathcal{D}(E^{(s)})] + \mathbb{E}_{\mathcal{D}^{(t)}} [\log (1 - \mathcal{D}(E^{(t)}))] \right\} \quad (7)$$

At the same time, the prediction network parameters are learned to maximize the domain classification loss (i.e., fooling the discriminator) on the target samples using the following optimization objective:

$$\min_{\theta_{\text{net}}} \mathbb{E}_{\mathcal{D}^{(t)}} [\log \mathcal{D}(E^{(t)})] \quad (8)$$

We use the hyperparameter λ_{adv} weighing the relative importance of the adversarial loss (8). Our training scheme

jointly optimizes the model parameters of the prediction network (θ_{net}) and the discriminator ($\theta_{\mathcal{D}}$). Updates to the prediction network and the discriminator happen upon every training iteration; however, when updating the prediction network, the discriminator parameters are kept fixed. Parameters of the discriminator are updated separately using the domain classification objective (Eq. 7).

3.5. Iterative Self Learning

Following prior work [71], we train our network end-to-end using an ISL scheme using Algorithm 1. We first train the prediction (θ_{net}) and discriminator ($\theta_{\mathcal{D}}$) networks for Q_1 iterations. We then generate semantic pseudo-labels ($\tilde{y}^{(t)}$) on the target training samples $x^{(t)}$ using the trained prediction network.

We further train the prediction network on the target training samples using pseudo-labels supervision and a masked cross-entropy loss (Eq. 1), masking target prediction pixels with confidence less than 0.9, for Q_3 iterations. Instead of training the prediction network using SL only once, we iterate over generating high-confidence pseudo-labels and self-training Q_2 times to refine the pseudo-labels, further resulting in better quality semantics output on the target domain.

We show in the ablation studies (Sec. 4.4) that our ISL scheme outperforms the simple SL. The discriminator network parameters ($\theta_{\mathcal{D}}$) are kept fixed during self-training.

Algorithm 1 $\text{ISL}(\mathcal{D}^{(s)}, \mathcal{D}^{(t)}, \theta_{\text{net}}, \theta_{\mathcal{D}})$

- 1: Train prediction (θ_{net}) and discriminator ($\theta_{\mathcal{D}}$) networks on source and target domains for Q_1 iterations;
 - 2: **for** Q_2 times **do**
 - 3: Generate $\tilde{y}_s^{(t)} = \mathcal{F}_s(\mathcal{F}_e(x^{(t)}))$ using trained θ_{net} ;
 - 4: Train θ_{net} on $(x^{(t)}, \tilde{y}_s^{(t)})$ for Q_3 iterations;
 - 5: **end for**
-

3.6. Network Architecture

The shared part of the prediction network $\mathcal{F}_e$ consists of a ResNet-101 backbone and a decoder (Fig. 2). The decoder consists of four convolutional layers; its outputs are fused with the backbone output features, which are denoted as the “shared feature map”. This shared feature map is then fed forward to the respective semantics and semantics refinement heads. Following the residual auxiliary block [41] (as in [59]), we place the depth prediction head between the last two convolutional layers of the decoder. In the supplementary materials, we show that our proposed approach is not sensitive to the residual auxiliary block and performs equally well with a standard multi-task learning network architecture (i.e., a shared encoder followed by multiple task-specific decoders). We adopt the Deeplab-V2 [4] architec-

tural design with Atrous Spatial Pyramid Pooling (ASPP) for the prediction heads. We use DC-GAN [47] as our domain discriminator for adversarial learning.

4. Experiments

4.1. UDA Benchmarks

We use three standard UDA evaluation protocols (EPs) to validate our model: **EP1**: SYNTHIA $\rightarrow$ Cityscapes (16 classes), **EP2**: SYNTHIA $\rightarrow$ Cityscapes (7 classes), and **EP3**: SYNTHIA $\rightarrow$ Mapillary (7 classes). A detailed explanation of these settings can be found in [59]. In all settings, the SYNTHIA dataset [49] is used as the synthetic source domain. In particular, we use the SYNTHIA-RAND-CITYSCAPES split consisting of 9,400 synthetic images and their corresponding pixel-wise semantic and depth annotations. For target domains, we use Cityscapes [10] and Mapillary Vistas [42] datasets. Following EP1, we train models on 16 classes common to SYNTHIA and Cityscapes; in EP2 and EP3, models are trained on 7 classes common to SYNTHIA, Cityscapes, and Mapillary. We use intersection-over-union to evaluate segmentation: IoU (class-IoU) and mIoU (mean-IoU). To promote reproducibility and emphasize significance of our results, we report two outcomes: the best mIoU, and the confidence interval. The latter is denoted as $mean \pm std$ collected over five runs, thus describing a 68% confidence interval centered at $mean^1$. For depth, we use Absolute Relative Difference ($|Rel|$), Squared Relative Difference (Rel^2), Root Mean Squared Error (RMS), its log-variant LRMS; and the accuracy metrics [14] as denoted by δ_1 , δ_2 , and δ_3 . For each metric, we use $\uparrow$ and $\downarrow$ to denote the improvement direction.

4.2. Experimental Setup

All our experiments are implemented in PyTorch [45]. Backbone network is a ResNet-101 [20] initialized with ImageNet [12] weights. The prediction and discriminator networks are optimized with SGD [2] and Adam [28] with learning rates 2.5×10^{-4} and 10^{-4} respectively. Throughout our experiments, we use $\lambda_r = 1.0$, $\lambda_d = \lambda_{adv} = 10^{-3}$. For generating depth bins, we use $Z_{\min} = 1\text{m}$, $Z_{\max} = 655.36\text{m}$, and $K = 15$. In all ISL experiments, parameters of the algorithm are: $Q_1 = 65K$, $Q_2 = 5$, $Q_3 = 5K$. Link to the project page with source code is in the Abstract.

4.3. Comparison to Prior Art

4.3.1 EP1

Table 1 reports semantic segmentation performance of our proposed model trained and evaluated following EP1. For a fair comparison with [53, 54, 39], we also report results

¹Class-IoU values of the “best mIoU” setting can be less than the mean of the class confidence interval at the expense of other classes performance.

Table 1: Semantic segmentation performance (IoU and mIoU, %) comparison to the prior art. All models are trained and evaluated using the EP1 protocol. mIoU* is computed on a subset of 13 classes, excluding those marked with *. For our method, we report the results of the run giving the best mIoU, as well as 68% confidence interval over five runs as *mean*±*std*.

		SYNTHIA → Cityscapes (16 classes)																mIoU ↑	mIoU* ↑
Models	Depth	road	sidewalk	building	wall*	fence*	pole*	light	sign	veg	sky	person	rider	car	bus	mbike	bike		
SPIGAN-no-PI [32]		69.5	29.4	68.7	4.4	0.3	32.4	5.8	15.0	81.0	78.7	52.2	13.1	72.8	23.6	7.9	18.7	35.8	41.2
SPIGAN [32]	✓	71.1	29.8	71.4	3.7	0.3	33.2	6.4	15.6	81.2	78.9	52.7	13.1	75.9	25.5	10.0	20.5	36.8	42.4
AdaptSegnet [53]		79.2	37.2	78.8	-	-	-	9.9	10.5	78.2	80.5	53.5	19.6	67.0	29.5	21.6	31.3	-	45.9
AdaptPatch [54]		82.2	39.4	79.4	-	-	-	6.5	10.8	77.8	82.0	54.9	21.1	67.7	30.7	17.8	32.2	-	46.3
CLAN [39]		81.3	37.0	80.1	-	-	-	16.1	13.7	78.2	81.5	53.4	21.2	73.0	32.9	22.6	30.7	-	47.8
Advent [58]		87.0	44.1	79.7	9.6	0.6	24.3	4.8	7.2	80.1	83.6	56.4	23.7	72.7	32.6	12.8	33.7	40.8	47.6
DADA [59]	✓	89.2	44.8	81.4	6.8	0.3	26.2	8.6	11.1	81.8	84.0	54.7	19.3	79.7	40.7	14.0	38.8	42.6	49.8
Ours (best mIoU)	✓	86.9	43.0	80.7	19.2	0.9	27.2	11.6	12.6	81.3	83.2	60.7	24.0	84.2	46.2	22.0	44.2	45.5	52.4
Ours (confidence)	✓	86.4 ±0.6	42.5 ±0.8	80.4 ±0.2	20.0 ±2.0	1.0 ±0.6	27.7 ±0.3	10.5 ±0.9	13.3 ±0.7	80.6 ±0.4	82.6 ±0.5	61.0 ±0.4	23.7 ±1.2	81.8 ±2.2	42.9 ±3.8	21.0 ±3.2	44.7 ±2.4	45.0 ±0.3	51.5 ±0.4

Table 2: Semantic segmentation performance (IoU and mIoU, %) comparison to the prior art. All models are trained and evaluated using the EP2 and EP3 protocols at different resolutions, as indicated in the resolution (“Res.”) column. For our method, we report the results of the run giving the best mIoU, as well as 68% confidence interval over five runs as *mean*±*std*.

		(a) SYNTHIA → Cityscapes (7 classes)								(b) SYNTHIA → Mapillary (7 classes)								mIoU ↑
Res.	Model	Depth	flat	const	object	nature	sky	human	vehicle	mIoU ↑	flat	const	object	nature	sky	human	vehicle	mIoU ↑
320 × 640	SPIGAN-no-PI [32]		90.3	58.2	6.8	35.8	69.0	9.5	52.1	46.0	53.0	30.8	3.6	14.6	53.0	5.8	26.9	26.8
	SPIGAN [32]	✓	91.2	66.4	9.6	56.8	71.5	17.7	60.3	53.4	74.1	47.1	6.8	43.3	83.7	11.2	42.2	44.1
	Advent [58]		86.3	72.7	12.0	70.4	81.2	29.8	62.9	59.4	82.7	51.8	18.4	67.8	79.5	22.7	54.9	54.0
	DADA [59]	✓	89.6	76.0	16.3	74.4	78.3	43.8	65.7	63.4	83.8	53.7	20.5	62.1	84.5	26.6	59.2	55.8
	Ours (best mIoU)	✓	90.8	77.5	15.7	77.1	82.9	45.3	68.6	65.4	86.6	57.4	19.7	73.0	87.5	45.1	68.1	62.5
	Ours (confidence)	✓	90.1 ±0.5	76.7 ±0.4	15.7 ±0.9	76.3 ±0.7	82.2 ±1.1	44.1 ±2.3	68.2 ±1.0	64.7 ±0.5	86.8 ±0.3	58.6 ±0.7	17.0 ±2.3	70.8 ±1.4	88.9 ±0.8	44.8 ±2.7	67.9 ±0.9	62.1 ±0.4
Full	Advent [58]		89.6	77.8	22.1	76.3	81.4	54.7	68.7	67.2	86.9	58.8	30.5	74.1	85.1	48.3	72.5	65.2
	DADA [59]	✓	92.3	78.3	25.0	75.5	82.2	58.7	72.4	69.2*	86.7	62.1	34.9	75.9	88.6	51.1	73.8	67.6
	Oracle (only-target)		97.6	87.9	46.0	87.9	88.8	69.1	88.6	80.8	95.0	84.2	54.8	87.7	97.2	70.2	87.5	82.4
	Ours (best mIoU)	✓	92.4	80.7	27.7	78.1	83.6	59.0	78.6	71.4	88.5	59.2	27.8	79.4	85.7	64.4	79.6	69.2
	Ours (confidence)	✓	92.2 ±0.3	80.8 ±0.1	27.0 ±0.9	78.6 ±0.8	84.9 ±1.2	54.5 ±3.2	78.2 ±1.3	70.8 ±0.4	88.4 ±0.1	58.6 ±0.7	29.0 ±0.8	79.8 ±0.4	85.0 ±0.9	63.2 ±1.3	79.0 ±0.4	69.0 ±0.1

* The correct mean of class IoU values in Table 2 of [59].

on 13 classes and the standard 16 classes settings. Our method achieves SOTA performance in EP1 on both 16 and 13 classes, outperforming [59, 32] by large margins. Now we can identify the major class-specific improvements of our method over the SOTA [59] DADA. The major gains come from the following classes – “wall” (+12.4%), “mo-

torbike” (+8%), “person” (+6%), “bicycle” (+5.4%) and “rider” (+4.7%). Moreover, our method shows consistent improvements on classes underrepresented in the target domain: “light” (+3%), “sign” (+1.5%), “bicycle” (+5.4%), and “motorbike” (+8%). Fig. 3 shows the results of the qualitative comparison of our method with DADA [59].

Figure 3: Qualitative semantic segmentation results with EP1: SYNTHIA $\rightarrow$ Cityscapes (16 classes). (a) Images from Cityscapes validation set; (b) ground truth annotations; (c) DADA [59] predictions; (d) our model predictions. Our method demonstrates notable improvements over [59] on “bus”, “person”, “motorbike”, and “bicycle” classes as highlighted using the yellow boxes.

Figure 4: Qualitative semantic segmentation results with EP3: SYNTHIA $\rightarrow$ Mapillary-Vista (7 classes). Top: images from Mapillary validation set; Middle: ground truth annotations; Bottom: our model predictions.

Note that our model delineates small objects like “human”, “bicycle”, and “motorbike” more accurately than DADA.

4.3.2 EP2 and EP3

Table 2 presents the semantic segmentation results in EP2 and EP3 benchmarks. The models are evaluated on the Cityscapes and Mapillary validation sets on their common 7 classes. We also train and evaluate our model on the 320×640 resolution to obtain a fair comparison with the reference low-resolution models. In a similar vein, the proposed method outperforms the prior works in EP2 and EP3 benchmarks for both full- and low-resolution (640×320) settings. We further show in Sec. 4.5 that our approach achieves state-of-the-art performance without ISL in EP2 and EP3 in both full- and low-resolution settings. The proposed CTRL coupled with SRH demonstrates consistent improvements over three challenging benchmarks by capitalizing on the inherent semantic and depth correlations. In

Table 3: Ablation study of our method from Sec. 4.4.

Conf	SemSup	DepSup	SRHSup	SemAdv	DepAdv	SRHAdv	SL	ISL	mIoU (%) $\uparrow$
<i>C1</i>	✓								30.7
<i>C2</i>	✓	✓							35.2
<i>C3</i>		✓	✓						33.7
<i>C4</i>	✓	✓	✓						33.1
<i>C5</i>	✓			✓					40.8
<i>C6</i>	✓	✓		✓	✓				40.2
<i>C7</i>		✓	✓		✓	✓			39.5
<i>C8</i>	✓	✓	✓	✓	✓	✓			42.1
<i>C9</i>	✓	✓	✓	✓	✓	✓	✓		44.1
<i>C10</i>	✓	✓		✓	✓			✓	42.8
<i>C11</i>	✓	✓	✓	✓	✓	✓		✓	45.5

EP2 and EP3, our models show noticeable improvements over the state-of-the-art [59] with mIoU gains of +2.2% (EP2-full-res), 2% (EP2-low-res), +1.6% (EP3-full-res), +6.7% (EP3-low-res). Despite the challenging domain gap between SYNTHIA and Mapillary, our model shows significant improvement (+6.7%) in a low-resolution setting, which suggests robustness to scale changes.

4.4. Ablation Studies

A comprehensive ablation study is reported in Table 3. We trained 11 different models, each having a different configuration; these are denoted as *C1*, ..., *C11*. We use the following shortcuts in Table 3 to represent different combinations of settings: “Sem” – semantic, “Dep” – depth, “Sup” – supervision, “Adv” – adversarial, and “Conf” – configuration. Configurations *C1* to *C4* denote supervised learning settings without any adversarial training. These models are trained on the SYNTHIA dataset and evaluated on Cityscapes validation set. Configurations from *C5* to *C7* denote different combinations of supervised and adversarial losses on the semantics, depth, and semantics refinement heads. *C8* is the proposed model with CTRL, but without ISL. *C9* to *C11* are models trained with SL or ISL with or without SRH. *C5* to *C11* follow EP1 protocol: SYNTHIA $\rightarrow$ Cityscapes UDA training and evaluation setting.

C1 is trained using semantics label supervision without any depth information or adversarial learning. By enabling parts of the model and training procedure, we observed the following tendencies: *C2* & *C3*: depth supervision (either direct or through SRH) improves performance; *C4*: however, adding SRH on top of the depth head in the supervised learning setting does not bring improvements; *C5*: effectiveness of entropy map domain alignment in semantics fea-

Table 4: Effectiveness of the joint feature space learned by our method (w/o ISL) for robust domain alignment. Performance in mIoU; legend for “Ours” as in Table 2.

Model	S→C (FR)	S→M (FR)	S→C (LR)	S→M (LR)
DADA [59]	69.2	67.6	63.4	55.8
Ours (best mIoU)	70.6	67.6	64.9	58.8
Ours (confidence)	69.5 ± 0.6	66.5 ± 0.6	64.3 ± 0.4	58.5 ± 0.4

Table 5: Improvement over the state-of-the-art [59] in monocular depth estimation. The models are trained following SYNTHIA → Cityscapes (16 classes) UDA setting w/o ISL and evaluated on the Cityscapes validation set.

Model	$ Rel \downarrow$	$Rel^2 \downarrow$	$RMS \downarrow$	$LRMS \downarrow$	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
DADA [59]	0.6	10.8	17.0	4.4	0.14	0.28	0.41
Ours	0.3	6.3	14.8	0.6	0.30	0.58	0.77

ture space [58]; $C6$ and $C7$: domain alignment in the depth or refined semantics feature spaces do not bring any further improvements; $C8$: a combination of depth and SRH with task-specific semantics improves the performance (i.e., our CTRL model); $C9$: SL brings further improvement but not as good as with our ISL training scheme; $C10$: emphasizes the improvement over $C6$ with ISL enabled; $C11$: positive contribution of the SRH towards improving the overall model performance. Finally, we achieve state-of-the-art segmentation results (mIoU 45.5%) by combining the proposed CTRL, SRH, and ISL (configuration $C11$).

4.5. Additional Experimental Analysis

4.5.1 Effectiveness of the Joint UDA Feature Space

This section analyzes the effectiveness of joint feature space learned by the CTRL for unsupervised domain alignment. We train and evaluate our CTRL model without ISL on two UDA benchmarks: (a) EP2: SYNTHIA to Cityscapes 7 classes ($S \rightarrow C$) and (b) EP3: SYNTHIA to Mapillary 7 classes ($S \rightarrow M$) in both full- and low-resolution (FR and LR) settings. In Table 4, we show the segmentation performance of our model on these four different benchmark settings and compare it against the state-of-the-art DADA model [59]. Out of four settings, the proposed CTRL model (w/o ISL) outperforms the DADA model with mIoU gains of +1.4%, +1.5%, and +3% in three benchmark settings attesting the effectiveness of the joint feature space learned by the proposed CTRL.

Besides, we train both DADA and our model with ISL

and notice improvements in both the models with mIoU 43.5% (DADA) and 45.5% (ours). The superior quality of the predictions of our model, when used as pseudo labels, provides better supervision to the target semantics; the same can be observed in both our quantitative (Tables 1 and 2) and qualitative results (Figs. 3 and 4).

4.5.2 Monocular Depth Estimation Results

In this section, we show that our model not only improves semantic segmentation but also learns a better representation for monocular depth estimation. This intriguing property is of great importance for multi-task learning. According to [41], paying too much attention to depth is detrimental to the segmentation performance. Following [41], DADA [59] uses depth as purely auxiliary supervision. We observed that depth predictions of [59] are noisy (also admitted by the authors), resulting in failure cases. We conjecture that a proper architectural design choice coupled with a robust multi-tasking feature representation (encoding task-specific and cross-task relationship) improves both semantics and depth. In Table 5, we report the depth estimation evaluation results on the Cityscapes validation set of our method and compare it against the DADA model [59]. Training and evaluation are done following the EP1 protocol: SYNTHIA → Cityscapes (16 classes). We use Cityscapes disparity maps as ground truth depth pseudo-labels for evaluation. Table 5 demonstrates a consistent improvement of depth predictions with our method over [59].

5. Conclusion

We proposed a novel approach to semantic segmentation and monocular depth estimation within a UDA context. The main highlights of this work are: (1) a Cross-Task Relation Layer (CTRL), which learns a joint feature space for domain alignment; the joint space encodes both task-specific features and cross-task dependencies shown to be useful for UDA; (2) a semantic refinement head (SRH) aids in learning task correlations; (3) a depth discretizing technique facilitates learning distinctive relationship between different semantic classes and depth levels; (4) a simple yet effective iterative self-learning (ISL) scheme further improves the model’s performance by capitalizing on the high confident predictions in the target domain. Our comprehensive experimental analysis demonstrates that the proposed method consistently outperforms prior works on three challenging UDA benchmarks by a large margin.

Acknowledgments. The authors gratefully acknowledge the support by armasuisse. We thank Amazon Activate for EC2 credits and the anonymous reviewers for the valuable feedback and time spent.

References

- [1] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(12):2481–2495, 2017. 2
- [2] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010*, pages 177–186. Springer, 2010. 5
- [3] David Bruggemann, Menelaos Kanakis, Stamatios Georgoulis, and Luc Van Gool. Automated search for resource-efficient branched multi-task networks. *Brit. Mach. Vis. Conf.*, 2020. 2
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017. 2, 5
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Eur. Conf. Comput. Vis.*, pages 801–818, 2018. 2
- [6] Po-Yi Chen, Alexander H Liu, Yen-Cheng Liu, and Yu-Chiang Frank Wang. Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2624–2632, 2019. 2
- [7] Yuhua Chen, Wen Li, Xiaoran Chen, and Luc Van Gool. Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1841–1850, 2019. 2
- [8] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *CVPR*, 2018. 2
- [9] Yuhua Chen, Wen Li, and Luc Van Gool. Road: Reality oriented adaptation for semantic segmentation of urban scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7892–7901, 2018. 2
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3213–3223, 2016. 5
- [11] Gabriela Csurka. A comprehensive survey on domain adaptation for visual applications. In *Domain adaptation in computer vision applications*, pages 1–35. Springer, 2017. 1
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 248–255. Ieee, 2009. 5
- [13] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Adv. Neural Inform. Process. Syst.*, 2014. 2
- [14] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Adv. Neural Inform. Process. Syst.*, pages 2366–2374, 2014. 5
- [15] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2002–2011, 2018. 2, 4
- [16] Brian Fulkerson, Andrea Vedaldi, and Stefano Soatto. Class segmentation and object localization with superpixel neighborhoods. In *Int. Conf. Comput. Vis.*, pages 670–677. IEEE, 2009. 2
- [17] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning (ICML)*, pages 1180–1189, 2015. 2
- [18] Timnit Gebru, Judy Hoffman, and Li Fei-Fei. Fine-grained recognition in the wild: A multi-task domain adaptation approach. 2017. 2
- [19] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-guided representation learning for self-supervised monocular depth. 2020. 2
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. 5
- [21] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018. 4
- [22] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. 2
- [23] Javed Iqbal and Mohsen Ali. Mlsl: Multi-level self-supervised learning for domain adaptation with spatially independent and semantically consistent labeling. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 1864–1873, 2020. 2, 3
- [24] Huaizu Jiang, Deqing Sun, Varun Jampani, Zhaoyang Lv, Erik Learned-Miller, and Jan Kautz. Sense: A shared encoder network for scene-flow estimation. In *Int. Conf. Comput. Vis.*, pages 3195–3204, 2019. 2
- [25] Jianbo Jiao, Ying Cao, Yibing Song, and Rynson Lau. Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss. In *Eur. Conf. Comput. Vis.*, pages 53–69, 2018. 2
- [26] Menelaos Kanakis, David Bruggemann, Suman Saha, Stamatios Georgoulis, Anton Obukhov, and Luc Van Gool. Reparameterizing convolutions for incremental multi-task learning without task interference. In *Eur. Conf. Comput. Vis.*, 2020. 2
- [27] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7482–7491, 2018. 1, 2, 3
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *Int. Conf. Learn. Represent.*, 2014. 5
- [29] L’ubor Ladický, Paul Sturges, Karteek Alahari, Chris Russell, and Philip HS Torr. What, where and how many? combining object detectors and crfs. In *Eur. Conf. Comput. Vis.*, pages 424–437. Springer, 2010. 2

- [30] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *2016 Fourth international conference on 3D vision (3DV)*, pages 239–248. IEEE, 2016. 2, 3
- [31] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [32] Kuan-Hui Lee, German Ros, Jie Li, and Adrien Gaidon. Spigan: Privileged adversarial learning from simulation. *Int. Conf. Learn. Represent.*, 2019. 1, 6
- [33] Bo Li, Chunhua Shen, Yuchao Dai, Anton Van Den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1119–1127, 2015. 2
- [34] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Int. Conf. Comput. Vis.*, 2017. 2
- [35] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(10):2024–2039, Oct. 2016. 2
- [36] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3431–3440, 2015. 2
- [37] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. In *International Conference on Machine Learning (ICML)*, pages 97–105, 2015. 2
- [38] Hao Lu, Lei Zhang, Zhiguo Cao, Wei Wei, Ke Xian, Chunhua Shen, and Anton van den Hengel. When unsupervised domain adaptation meets tensor representations. In *Int. Conf. Comput. Vis.*, 2017. 2
- [39] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2507–2516, 2019. 5, 6
- [40] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive single-tasking of multiple tasks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1851–1860, 2019. 1, 2
- [41] Taylor Mordan, Nicolas Thome, Gilles Henaff, and Matthieu Cord. Revisiting multi-task learning with rock: a deep residual auxiliary block for visual detection. In *Adv. Neural Inform. Process. Syst.*, pages 1310–1322, 2018. 5, 8
- [42] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *Int. Conf. Comput. Vis.*, pages 4990–4999, 2017. 5
- [43] Davy Neven, Bert De Brabandere, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Fast scene understanding for autonomous driving. *arXiv preprint arXiv:1708.02550*, 2017. 2
- [44] Art B Owen. A robust hybrid of lasso and ridge regression. *Contemporary Mathematics*, 443(7):59–72, 2007. 2
- [45] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch, 2017. 5
- [46] Sujoy Paul, Yi-Hsuan Tsai, Samuel Schuster, Amit K Roy-Chowdhury, and Manmohan Chandraker. Domain adaptive semantic segmentation using weak labels. *arXiv preprint arXiv:2007.15176*, 2020. 2
- [47] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 5
- [48] Pierluigi Zama Ramirez, Matteo Poggi, Fabio Tosi, Stefano Mattoccia, and Luigi Di Stefano. Geometry meets semantics for semi-supervised monocular depth estimation. In *ACCV*, pages 298–313. Springer, 2018. 2
- [49] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3234–3243, 2016. 5
- [50] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Unsupervised domain adaptation for semantic segmentation with gans. *arXiv preprint arXiv:1711.06969*, 2:2, 2017. 2
- [51] Jamie Shotton, Matthew Johnson, and Roberto Cipolla. Semantic texon forests for image categorization and segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1–8. IEEE, 2008. 2
- [52] Trevor Standley, Amir R Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? *International Conference on Machine Learning (ICML)*, 2019. 2
- [53] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7472–7481, 2018. 2, 5, 6
- [54] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schuster, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *Int. Conf. Comput. Vis.*, pages 1456–1465, 2019. 5, 6
- [55] Simon Vandenhende, Stamatios Georgoulis, Bert De Brabandere, and Luc Van Gool. Branched multi-task networks: deciding what layers to share. *arXiv preprint arXiv:1904.02920*, 2019. 2
- [56] Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, Dengxin Dai, and Luc Van Gool. Revisiting multi-task learning in the deep learning era. *arXiv preprint arXiv:2004.13379*, 2020. 2
- [57] Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Mti-net: Multi-scale task interaction networks for multi-task learning. *Eur. Conf. Comput. Vis.*, 2020. 1, 2
- [58] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In

- IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2517–2526, 2019. [2](#), [4](#), [6](#), [8](#)
- [59] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Dada: Depth-aware domain adaptation in semantic segmentation. In *Int. Conf. Comput. Vis.*, pages 7364–7373, 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [60] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 675–684, 2018. [2](#)
- [61] Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. [2](#)
- [62] Jinyu Yang, Weizhi An, Sheng Wang, Xinliang Zhu, Chaochao Yan, and Junzhou Huang. Label-driven reconstruction for domain adaptation in semantic segmentation. *arXiv preprint arXiv:2003.04614*, 2020. [2](#)
- [63] Jinyu Yang, Weizhi An, Chaochao Yan, Peilin Zhao, and Junzhou Huang. Context-aware domain adaptation in semantic segmentation. *arXiv preprint arXiv:2003.04010*, 2020. [2](#)
- [64] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. 2016. [2](#)
- [65] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3712–3722, 2018. [2](#)
- [66] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. 2017. [2](#)
- [67] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Zequn Jie, Xiang Li, and Jian Yang. Joint task-recursive learning for semantic segmentation and depth estimation. In *Eur. Conf. Comput. Vis.*, pages 235–251, 2018. [2](#)
- [68] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Yan Yan, Nicu Sebe, and Jian Yang. Pattern-affinitive propagation across depth, surface normal and semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4106–4115, 2019. [2](#)
- [69] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2881–2890, 2017. [2](#)
- [70] Qianyu Zhou, Zhengyang Feng, Guangliang Cheng, Xin Tan, Jianping Shi, and Lizhuang Ma. Uncertainty-aware consistency regularization for cross-domain semantic segmentation. *arXiv preprint arXiv:2004.08878*, 2020. [2](#)
- [71] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Eur. Conf. Comput. Vis.*, pages 289–305, 2018. [2](#), [5](#)
- [72] Laurent Zwald and Sophie Lambert-Lacroix. The berhu penalty and the grouped effect. *arXiv preprint arXiv:1207.6868*, 2012. [2](#)