

DADA: Depth-Aware Domain Adaptation in Semantic Segmentation

Tuan-Hung Vu¹ Himalaya Jain¹ Maxime Bucher¹ Matthieu Cord^{1,2} Patrick Pérez¹

¹valeo.ai, Paris, France

²Sorbonne University, Paris, France

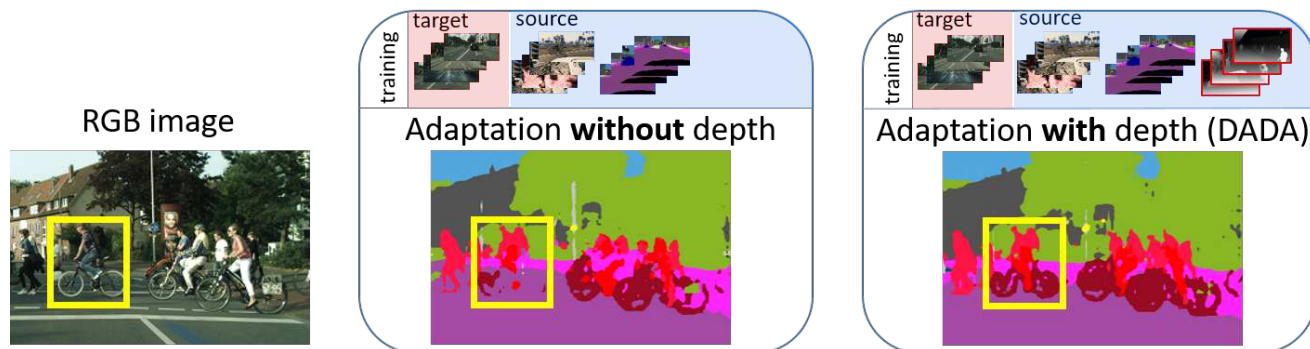


Figure 1: We propose a novel depth-aware domain adaptation framework (DADA) to efficiently leverage depth as privileged information in the unsupervised domain adaptation setting. This example shows how semantic segmentation of a scene from the target domain benefits from the proposed approach, in comparison to state-of-the-art domain adaptation with no use of depth. In figure’s top, we use different background colors (blue and red) to represent source and target information that are available during training. Here, annotated source domain data come from the synthetic SYNTHIA dataset and un-annotated target domain images are real scenes from Cityscapes. The cyclist highlighted by the yellow box is a good qualitative illustration of the improvement we obtain.

Abstract

Unsupervised domain adaptation (UDA) is important for applications where large scale annotation of representative data is challenging. For semantic segmentation in particular, it helps deploy, on real “target domain” data, models that are trained on annotated images from a different “source domain”, notably a virtual environment. To this end, most previous works consider semantic segmentation as the only mode of supervision for source domain data, while ignoring other, possibly available, information like depth. In this work, we aim at exploiting at best such a privileged information while training the UDA model. We propose a unified depth-aware UDA framework that leverages in several complementary ways the knowledge of dense depth in the source domain. As a result, the performance of the trained semantic segmentation model on the target domain is boosted. Our novel approach indeed achieves state-of-the-art performance on different challenging synthetic-2-real benchmarks. Code and models are available at <https://github.com/valeoai/DADA>.

1. Introduction

Advances in deep convolutional neural networks (CNNs) brought significant leaps forward in many recognition tasks including semantic segmentation. Still, predicting semantic labels for all imagery pixels is a challenging problem, especially when models are trained on one domain, named as *source*, yet evaluated on another domain, named as *target*. The so-called *domain gap* between source and target distributions often causes drastic drops in target performance. Instead, autonomous critical systems such as self-driving cars require robust performance under diverse testing conditions, despite the lack of ubiquitous training data. For such systems, semantic segmentation models trained on sunny urban images taken in Roma should yield good results even on foggy scenes in London.

Techniques addressing the domain gap problem are usually classified as *domain adaptation* (DA) [7]. In previous works, most DA settings are *unsupervised* on the target side, *i.e.*, only un-annotated target samples are available during the supervised training on source domain. This is referred to as *unsupervised domain adaptation* (UDA). In recent years, considerable progress has been made in UDA

with approaches such as distribution discrepancy minimization with MMD [20, 42] or adversarial training [14, 35], generative approaches [13] and others.

In this UDA context, the *synthetic-2-real* scenario, where source and target samples are synthetic and real images respectively, is especially appealing thanks to the zero-cost source label acquisition. Moreover, recent virtual engines can simulate other sensory outputs like dense/sparse depth or radar; for example, the SYNTHIA dataset [29] provides corresponding virtual depth maps of synthesized urban scenes. Most previous UDA works ignore such extra information, except for Lee *et al.* [18] who propose to use depth for regularizing a style-transfer network.

In this work, we propose a new scheme to leverage depth information available in source domain for UDA. We hypothesize that introducing additional depth-specific adaptation brings complementary effects to further bridge the performance gap between source and target at test time. Toward this end, we transform the segmentation backbone such that the depth information is embedded into a dedicated deep architecture by means of an auxiliary depth regression task. Depth, operating as an additional source-domain supervision in our framework (only available while training), will be considered as a *privileged information*.

Another challenge is to incorporate efficiently depth signals into the UDA learning. Addressing this concern, we introduce a new depth-aware adversarial training protocol based on the fusion of the network outputs. Such a late fusion was inspired by our intuition that visual information at different depth levels should be treated differently.

The proposed approach is illustrated in Figure 1, where the benefit of depth-aware adaptation on key object categories like ‘human’ and ‘vehicle’ is visible. The contributions of this approach are the following:

- Depth-aware UDA learning strategy: we introduce a novel depth-aware adaptation scheme, coined DADA learning, which simultaneously aligns segmentation-based and depth-based information of source and target while being aware of scene geometry.
- Depth-aware deep architecture: we propose a novel depth-aware segmentation pipeline, named DADA architecture, in which depth-specific and standard CNN appearance features are fused before being forwarded through the segmentation classifiers.
- State-of-the-art performance: evaluations show that our framework achieves SotA results on challenging *synthetic-2-real* benchmarks. We also report ablation studies to provide insights into the proposed approach.

2. Related works

Unsupervised domain adaption has received a lot of attention in last few years [7, 9, 14, 21, 37, 42]. It is usually

approached with domain discrepancy minimization, generative approaches or using some privileged information to guide the learning. Since we are only concerned with visual semantic segmentation in this work, we limit our review of UDA to approaches that aim at this task as well.

Various approaches of UDA for segmentation employ adversarial training to minimize cross-domain discrepancy. The main idea, stemming from generative adversarial networks (GANs) [10], is to train a discriminator for predicting the domain of the data (source or target) while the segmentation network tries to fool it (along with the supervised segmentation task on the source). Under the competition with the discriminator, the segmentation network tries to map its input to domain-agnostic intermediate or final representation (which is the input to the discriminator), before accomplishing its task. This alignment with adversarial training is usually done in the feature space. In [14], the feature alignment is done not only with adversarial training but also by transferring the label statistics of the source domain by category specific adaptation. [5] uses adversarial training for class-level alignment on grid-wise soft pseudo-labels. In [4], spatial-aware adaptation is conducted and a distillation loss is used to address specifically synthetic-to-real adaptation by enforcing the segmentation network’s output to be similar to a reference network trained on real data. While most of these methods do alignment on the features space, recently [35, 39] propose alignment on the output space. [35] does the alignment on the prediction of the segmentation network and [39] proposes to do it on the *weighted self-information* of the prediction probability. [36] and [23] extend the approach of [35] by patch-level alignment and category-level adversarial loss respectively. Another use of adversarial training for UDA is proposed in [30, 31], where the discrepancy between two instances of the same input from target domain is minimized while the classification layer tries to maximize it. The approach we propose also uses adversarial training but it takes this further by introducing the depth.

Another strategy that received much attention is the use of generative networks to turn source domain samples into target-like images. CyCADA [13] uses Cycle-GAN [45] to generate target-like images conditioned on the source images, *i.e.*, generated images contain the structure or semantic content of the source with the “style” (colors and textures) of the target domain. These generated images inherit the ground truth semantic segmentation of the conditioning source images and can then be used for supervised training of the segmentation network. [46] aligns the source and target embeddings using GAN [10] and replaces the cross-entropy loss by a *conservative loss* (CL) which penalizes the easy and hard source examples. In DCAN [41], a similar generative approach is used with channel-wise feature alignment in the generator and segmentation networks.

Some other interesting works have investigated combination of adversarial and generative approaches [25, 32, 44], self-training [47] and curriculum style learning [43].

In order to help domain adaptation in various tasks, several works explore the use of *privileged information* on the source data [2, 19, 33]. The idea of using privileged information (PI) for learning was first formulated by Vapnik & Vashist (2009) [38]. PI is an additional information available only at training time. This is conceptually similar to humans learning new notions or concepts with the help of teacher’s comment or explanation. Following [38] many works [12, 22, 24, 34, 40] investigated PI for various tasks. Recently, SPIGAN [18] has proposed such a UDA approach for semantic segmentation. SPIGAN first uses a generative network to convert the source images into target-type images. These new images are then used to train the segmentation network along with a depth regression network (privileged information network) in a supervised manner, thanks to the ground truth of the source images. This work shows that the additional depth regression task helps in pixel level adaptation or, in other words, better captures the content of source images in the generated target-like images.

In the present work, we also leverage a depth regression task to aid domain adaptation. However, this is accomplished in a way that notably differs from SPIGAN (which uses depth only as a regularization for the generator).

3. Depth-aware domain adaptation

In this section, we describe our proposed UDA method for semantic segmentation using depth. Our goal is to use depth as the privileged information in the UDA setting to improve segmentation performance on the target domain.

To this end, we modify a network of semantic segmentation by including a monocular depth regression part. More specifically, we design a deep architecture to embed the depth in a dedicated residual block as illustrated in Figure 2 (‘DADA architecture’ in top part). In Section 3.1, we detail DADA network architecture and the supervised learning used on source dataset. The second part concerns the learning scheme for such a UDA approach. To get the full benefit of the geometry information, we propose a depth-aware adversarial learning scheme. We argue that for domain adaptation, particularly in the urban settings, objects closer to the ego-camera should be emphasized more. Our framework, illustrated in Figure 2 (‘DADA learning’ in bottom part), is detailed in Section 3.2.

3.1. DADA Network Architecture.

Starting from an existing semantic segmentation architecture, we insert additional modules (1) to predict monocular depth as additional output and (2) to feed the information exploited by this auxiliary task back to the main stream. More specifically, we adopt the *residual auxiliary*

block that was recently introduced in [24] for detection. It amounts to grafting a new branch to the backbone CNN. In this branch, the backbone CNN features are consecutively fed into three encoding convolutional layers, followed by an average pooling layer to output depth map predictions. On the residual path back to the main branch, the encoded features (before the depth pooling) are decoded by a convolutional layer and fused with the backbone features. The top-part of Figure 2 shows the proposed hybrid architecture, mixing the auxiliary block architecture with the backbone one.

Importantly, for the feature-level fusion, we adopt an element-wise product, indicated as “Feat. Fusion” in the top part of Figure 2. To produce segmentation predictions, we feed-forward the fused features through the remaining classification modules.

Source domain supervised training. Our model is trained with supervised segmentation and depth losses on source domain. We consider a training set $\mathcal{T}_s \subset \mathbb{R}^{H \times W \times 3} \times (1, C)^{H \times W} \times \mathbb{R}_+^{H \times W}$ of source color images of size $H \times W$ along with pixel-level C -class segmentation and depth annotations. Let $DADA$ be the network which takes an image \mathbf{x} and jointly predicts a C -dimensional “soft-segmentation map” $DADA_{\text{seg}}(\mathbf{x}) = \mathbf{P}_{\mathbf{x}} = [\mathbf{P}_{\mathbf{x}}^{(h,w,c)}]_{h,w,c}^1$ and a depth map $DADA_{\text{depth}}(\mathbf{x}) = [\mathbf{Z}_{\mathbf{x}}^{(h,w)}]_{h,w}$. Similar to [15], we adopt the inverse depth representation, *i.e.*, depth attenuates when moving away from the camera. The parameters θ_{DADA} of $DADA$ are learned to minimize the segmentation and depth losses on source samples $(\mathbf{x}_s, \mathbf{y}_s, \mathbf{z}_s) \in \mathcal{T}_s$:

$$\mathcal{L}_{\text{seg}}(\mathbf{x}_s, \mathbf{y}_s) = - \sum_{h=1}^H \sum_{w=1}^W \sum_{c=1}^C \mathbf{y}_s^{(h,w,c)} \log \mathbf{P}_{\mathbf{x}_s}^{(h,w,c)}, \quad (1)$$

$$\mathcal{L}_{\text{dep}}(\mathbf{x}_s, \mathbf{z}_s) = - \sum_{h=1}^H \sum_{w=1}^W \text{berHu}(\mathbf{Z}_{\mathbf{x}_s}^{(h,w)} - \mathbf{z}_s^{(h,w)}), \quad (2)$$

with the reverse Huber loss defined as [17]:

$$\text{berHu}(e_z) = \begin{cases} |e_z|, & \text{if } |e_z| \leq c, \\ \frac{e_z^2 + c^2}{2c}, & \text{otherwise,} \end{cases} \quad (3)$$

where c is a positive threshold that we fix in practice to $\frac{1}{5}$ of the maximum depth residual. Empirically, the berHu loss is favorable for the depth regression task: samples with larger residuals are penalized more by the ℓ_2 term, while gradients of small-residual samples are more underlined with ℓ_1 .

Finally, our DADA optimization problem on source domain is formulated as:

$$\min_{\theta_{DADA}} \frac{1}{|\mathcal{T}_s|} \sum_{\mathcal{T}_s} \mathcal{L}_{\text{seg}}(\mathbf{x}_s, \mathbf{y}_s) + \lambda_{\text{dep}} \mathcal{L}_{\text{dep}}(\mathbf{x}_s, \mathbf{z}_s), \quad (4)$$

with λ_{dep} a weighting factor for depth regression.

¹Seen as empirical probabilities that sum to one over the C classes.

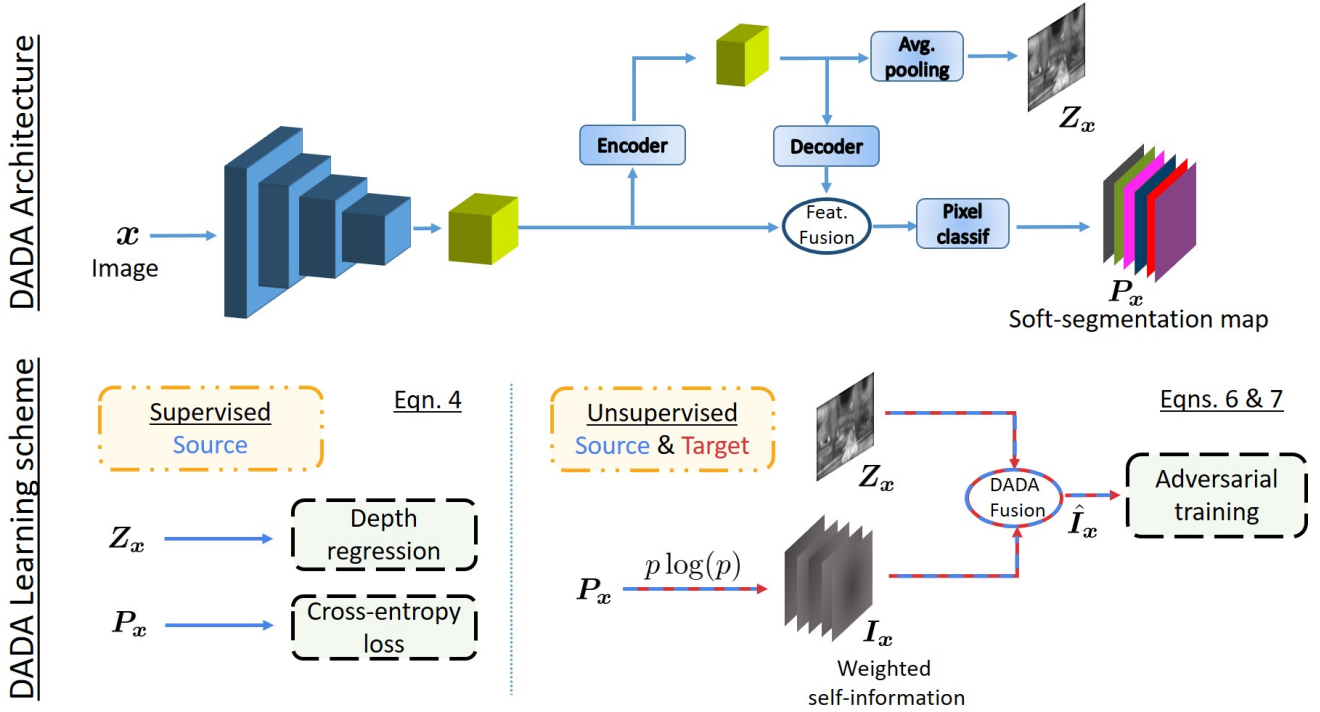


Figure 2: **DADA architecture (top) and DADA learning scheme (bottom)**. In the top part, the dark-blue stack shows the backbone CNN network; light-blue boxes symbolize the network modules; and green blocks stand for output features. In the lower part, the arrows drawn in blue and red differentiate network flows of source and target samples respectively. For convenient reference, over the learning blocks – illustrated by dashed boxes – we indicate the corresponding equation numbers.

3.2. DADA Adversarial Learning Scheme

For UDA in semantic segmentation, the key idea is to align the source and target domains so that a discriminator network cannot distinguish between the domains. We follow here the recent strategies that align features at the output level [35, 39], *i.e.*, the soft segmentation map P_x produced by the segmentation network $DADA_{seg}$ on input image x .

We hereby question the plausibility of adapting such a methodology for an auxiliary space, *i.e.*, depth prediction Z_x in the present work, in the hope that the performance of the main task is improved. We hypothesize that aligning source and target distributions also in the depth space implicitly bridges the domain gaps of the shared lower-level CNN representations and should bring improvements to the main task on target domain.

In order to carry out such a strategy, we propose a joint alignment: we first merge both signals, then we feed-forward the fused features as input to a discriminator.

We illustrate DADA adversarial learning scheme in the lower right part of Figure 2. More precisely:

- We compute weighted self-information (“surprisal”) maps [39] $I_x \in [0, 1]^{H \times W \times C}$ defined as:

$$I_x^{(h,w,c)} = -P_x^{(h,w,c)} \cdot \log P_x^{(h,w,c)}. \quad (5)$$

- We fuse the weighted self-information I_x with the depth prediction Z_x to produce a depth-aware map \hat{I}_x . The fusion of Z_x and I_x , which we refer to as *DADA fusion*, is the element-wise product of I_x with Z_x . As here the inverse depth is used, such a multiplication implies stronger attention toward scene elements that are closer to the ego-camera.
- Then, we do the adversarial adaptation on \hat{I}_x .

The depth-aware map \hat{I}_x carries 3D-structural and geometrical information which should be consistent across the domains, thus alignment on this space is beneficial to adaptation as we will see in the next section.

Formally, given \mathcal{X}_t the set of un-annotated images in the target domain, the discriminator D is trained to distinguish source *vs.* target outputs (labeled as ‘1’ and ‘0’ respectively) with the following classification objective minimization:

$$\min_{\theta_D} \frac{1}{|\mathcal{T}_s|} \sum_{\mathcal{T}_s} \mathcal{L}_D(\hat{I}_{x_s}, 1) + \frac{1}{|\mathcal{X}_t|} \sum_{\mathcal{X}_t} \mathcal{L}_D(\hat{I}_{x_t}, 0), \quad (6)$$

and the DADA network is updated using the “fooling” objective minimization:

$$\min_{\theta_{DADA}} \frac{1}{|\mathcal{X}_t|} \sum_{\mathcal{X}_t} \mathcal{L}_D(\hat{I}_{x_t}, 1). \quad (7)$$

At each training iteration, we feed the network a mini-batch of two samples coming from source and target domains. The two objectives in (4) and (7) are jointly optimized with a weighting factor λ_{adv} used for the adversarial part. Gradients of all losses are accumulated and then back-propagated to update the network.

Discussion. The only method using depth as privileged information in the same task as ours is SPIGAN [18]. While in DADA, with feature fusion and DADA fusion, we exploit depth to enhance the appearance features and improve the source-target alignment, SPIGAN leverages depth as a regularization for the pixel-level alignment generator. We argue that our way of using depth is more explicit, which, in return, gets more benefit from the privileged information. Moreover, thanks to the residual fusion in the auxiliary block presented in Section 3.1, depth signals are directly taken into account while deriving main task predictions. Such a residual fusion spreads beneficial effects of the depth-specific adaptation to the main segmentation task.

4. Experiments

This section presents quantitative and qualitative results. We introduce in Section 4.1 the *synthetic-2-real* benchmarks used in this work. We then analyze the performance of the proposed model in Section 4.2 and report ablation studies in Section 4.3.

4.1. Experimental details

Datasets. In this work, we use SYNTHIA dataset [29] as the source domain. It is composed of 9,400 synthetic images annotated with pixel-wise semantic labels and depth. Similar to previous works [18, 35, 39], we adopt the split SYNTHIA-RAND-CITYSCAPES using Cityscapes-style annotations. For the target domain, we use either Cityscapes [6] or Mapillary Vistas [26] datasets. What follows are our experimental set-ups in detail:

- *SYNTHIA*→*Cityscapes* (16 classes): This is a standard evaluation protocol used in previous works. The models are trained on the 16 classes common to SYNTHIA and Cityscapes. Similar to [35, 39], we also report performance on the 13-class subset.
- *SYNTHIA*→*Cityscapes* / *Vistas* (7 classes): Following [18], we conduct experiments on the 7 categories that are common to SYNTHIA, Cityscapes and Vistas.

Network architecture. In our experiments, we adopt Deeplab-V2 [3] based on ResNet-101 [11] as the backbone segmentation architecture. Like [35, 39], we apply Atrous Spatial Pyramid Pooling (ASPP) with sampling rates of {6, 12, 18, 24}. Segmentation prediction is only done on the

conv5 features. For the adversarial training, we use DCGAN’s discriminator [28] composed of 4 sequential convolutional layers with leaky-ReLUs as activation functions.

The encoding module used for depth regression has three consecutive convolutional layers: the first and last ones have kernel size of 1; the middle layer has kernel size of 3 with a suitable zero-padding to ensure the same input and output resolutions. Each layer uses 4 times fewer channels than the previous one. In the decoding part, we feed the encoded features through a 1×1 convolutional layer. The decoding layer has the same number of output channels as the channel size of the ResNet-101 backbone feature.

Implementation details. Implementations are done with the PyTorch deep learning framework [27]. To train and validate our models, we use a single NVIDIA 1080TI GPU with 11GB memory. We initialize our models with the ResNet-101 [11] pre-trained on the ImageNet dataset [8]. Segmentation and depth regression networks are trained by a standard Stochastic Gradient Descent optimizer [1] with learning rate 2.5×10^{-4} , momentum 0.9 and weight decay 10^{-4} . For discriminator training, we adopt Adam optimizer [16] with learning rate 10^{-4} . In all experiments, we fixed λ_{dep} as 10^{-3} for the depth regression task and used $\lambda_{\text{adv}} = 10^{-3}$ to weight the adversarial loss.

4.2. Results

We present results of the proposed DADA approach in comparison to different baselines. On the three benchmarks, our models achieve state-of-the-art performance. Our extensive study shows the benefit of leveraging depth as privileged information with our DADA framework for UDA in semantic segmentation.

SYNTHIA→**Cityscapes:** In Table 1, we report semantic segmentation performance in term of “mean Intersection over Union” (mIoU in %) on the 16 classes of the Cityscapes validation set. DADA achieves state-of-the-art performance on the benchmark. To the best of our knowledge, SPIGAN [18] is the only published work targeting the same problem that also considers depth as privileged information. DADA achieves a D-Gain of 1.8%, almost double of SPIGAN’s. Analyzing per-class results, we observe that the improvement over AdvEnt [39] primarily comes from the ‘vehicle’ category, *i.e.*, ‘car’ (+7%), ‘bus’ (+8.1%) and ‘bike’ (+5.1%). On ‘object’ classes like ‘light’ and ‘pole’, DADA introduces moderate gains. Figure 3 illustrates some qualitative examples comparing DADA and the AdvEnt baseline. Our model shows better results on ‘vehicle’ classes while AdvEnt sometimes makes severe mistakes of predicting ‘car’ as ‘road’ or ‘sidewalk’. DADA also outperforms significantly other baseline methods that report results on a 13-class subset.

SYNTHIA → Cityscapes (16 classes)

Models	Depth	road	sidewalk	building	wall*	fence*	pole*	light	sign	veg	sky	person	rider	car	bus	mbike	bike	mIoU	D-Gain	mIoU*
SPIGAN-no-PI [18]		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 [18]	✓	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	1.0	42.4
AdaptSegnet [35]		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 [36]		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 [23]		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 [39]		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	✓	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	1.8	49.8

Table 1: **Semantic segmentation performance mIoU (%) on Cityscapes validation set of different models trained on SYNTHIA.** Top and bottom sub-tables correspond to VGG-16-based and ResNet-101-based models respectively. For methods making use of depth, we report the absolute depth-driven mIoU gain (D-Gain). We also show the mIoU (%) of the 13 classes (mIoU*) excluding classes with *.

		(a) SYNTHIA → Cityscapes (7 classes)										(b) SYNTHIA → Vistas (7 classes)							
Reso.	Models	Depth	flat	const.	object	nature	sky	human	vehicle	mIoU	flat	const.	object	nature	sky	human	vehicle	mIoU	
320 × 640	SPIGAN-no-PI [18]		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 [18]	✓	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* [39]		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	✓	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	
Full	Advent* [39]		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	✓	92.3	78.3	25.0	75.5	82.2	58.7	72.4	70.4	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	

Table 2: **Semantic segmentation performance mIoU (%) in 7-classes setups.** (a) Cityscapes and (b) Vistas validation set. We report results produced at different resolutions. AdvEnt* is the adaptation of AdvEnt published code to 7-classes set-ups.

Table 2-a shows results in the same experimental set-up, except that training and validation were done on 7 categories. To be comparable with [18], we report additional results produced at the 320×640 resolution. Our DADA framework outperforms the state-of-the-art on this benchmark by a large margin. Over the AdvEnt baseline, we achieve +3.2% mIoU improvement. Similar to the 16-class results, an important gain on the ‘vehicle’ category (+3.7%) is observed. In addition, we contrast a +3% IoU on the ‘human’ category to the negative IoU drops on classes ‘person’ and ‘rider’ reported in Table 1. We conjecture that this drop stems from the intra-category confusion, *i.e.*, ‘pedestrian’ and ‘rider’ are easily confused. A significant improvement is pointed out in the lower resolution set-up where using depth adds +14.0% to the ‘human’ category IoU. These re-

sults demonstrate the merit of DADA for UDA, especially on crucial categories like ‘human’ and ‘vehicle’ – the vulnerable road users (VRUs).

An interesting UDA metric introduced in [18] is the negative transfer rate (lower is better) – the percentage of after-adaptation test cases having per-image mIoUs lower than ones coming from the model trained only on source (without adaptation). On SYNTHIA→Cityscapes (7 classes), 320×640 resolution DADA model has only 5% negative transfer rate, compared to 9% for SPIGAN. It is worth noting that our only-on-source mIoU in this case is 50%, much larger than the one reported in [18] (36.3%). The AdvEnt baseline, without using depth, suffers from a negative transfer rate of 11% – more than double of DADA’s.




Figure 3: **Qualitative results in the SYNTHIA→Cityscapes (16 classes) set-up.** The four columns plot (a) RGB input images, (b) ground-truths, (c) AdvEnt baseline outputs and (d) DADA predictions. DADA shows good performance on ‘bus’, ‘car’, ‘bicycle’ classes. Best viewed in color.

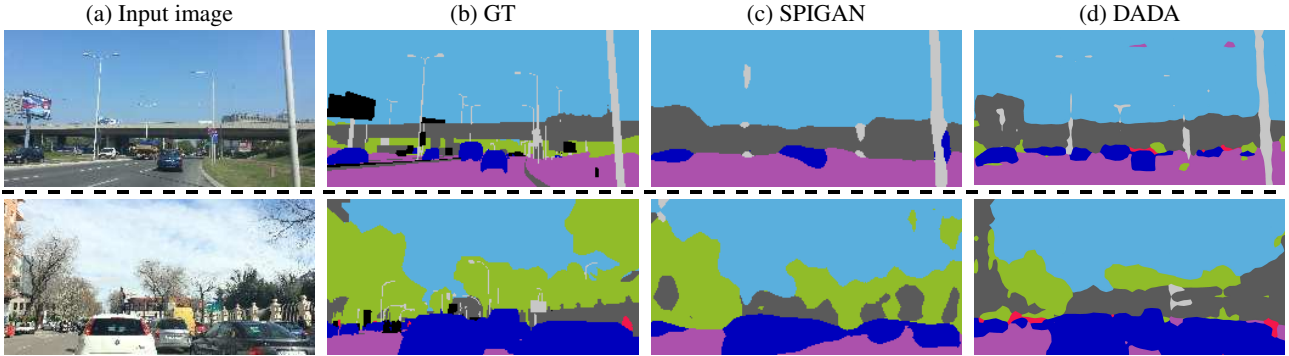


Figure 4: **Qualitative results in the SYNTHIA→Vistas (7 classes) set-up.** All models were trained and tested at the resolution of 320×640 . From left to right, we show (a) RGB input images, (b) corresponding segmentation ground-truths, (c) SPIGAN’s and (d) our DADA’s segmentation predictions. Not only DADA performs visually better than SPIGAN, but it also produces correct predictions on wrongly annotated construction areas. Best viewed in color.

SYNTHIA→Vistas: In this experiment, the Mapillary Vistas [26] is used as the target domain. In SPIGAN [18], the authors report unfavorable UDA behaviors on Vistas compared to Cityscapes. This seems caused by the artifacts that the source-target image translation introduces when trying to close the larger gap between SYNTHIA and Vistas. In such a case, leveraging depth information demonstrates important adaptation improvement (+17.3%). On the other hand, our UDA framework does not undergo such a difficulty. Indeed, as shown in Table 2(a-b), the AdvEnt baseline

performs much better than SPIGAN-no-PI, with no significant difference in absolute mIoU on the two target datasets (59.4% vs. 54.0%). Over such a stronger baseline, DADA still achieves an overall improvement of +2.4% mIoU. We also obtain best per-class IoUs on the benchmark.

DADA (trained and tested on 320×640 images) has 30% negative transfer rate compared to the 42% of SPIGAN. As discussed in [18], the challenging domain gap between SYNTHIA and Vistas might cause these high negative rates. In addition to this explanation, we also ques-

Setup	Surp. Adapt.	Depth Adapt.	Feat. Fusion	DADA Fusion	mIoU (%)
S1 (no adapt.)					32.2
S2 (AdvEnt)	✓				40.8
S3	✓		✓		40.7
S4		✓			35.7
S5		✓	✓		38.0
S6	✓	✓	✓		41.6
S7 (DADA)	✓	✓	✓	✓	42.6

Table 3: **Segmentation performance (mIoU) on the Cityscapes validation set of 7 ablation experiments.** Setup S1, with no check-marks, indicates source-only training.

tion the annotation quality of the Vistas dataset, visually inspection of results having revealed inconsistencies. Interestingly, when we evaluate the DADA model trained with the current set-up (SYNTHIA→Vistas) on the Cityscapes validation set with arguably cleaner annotations, the obtained negative transfer rate reduces to 6%.

In Figure 4, we show some qualitative results comparing our best model with SPIGAN. As mentioned above, we note that the Vistas segmentation annotations are noisy. For example, some construction areas slightly covered by tree branches are annotated as ‘vegetation’. DADA provides reasonable predictions on these areas – sometimes even better than human ground-truths.

4.3. Ablation studies

Effect of depth-aware adversarial adaptation. We report in Table 3 performance with seven training setups, S1 to S7: S1 is the source-only baseline (no adaption at all), S2 amounts to AdvEnt (no use of depth) and S7 is DADA. Intermediate setups S3 to S6 amount to using or not the AdvEnt’s “surprisal adaptation”, the auxiliary depth adaptation (“depth adaptation”), the point-wise feature fusion mechanism (“feature fusion”) and the output-based fusion for DADA adversarial training (“DADA fusion”). First, we remark that adversarial adaptation on the auxiliary depth-space (S4 and S5) does help improve performance of the main task. Improvement of S5 over S4 demonstrates the advantage of depth fusion at the feature-level. Comparable performance of S2 and S3 indicates that depth supervision on the source domain is not effective in absence of depth-specific adaptation. Indeed, S6, with two separate adversarial adaptations on the surprisal and depth spaces, works better than S2 and S3. Still, in S6, the coupling between spaces

% of SYNTHIA	10%	30%	50%	70%	100%
Cityscapes mIoU	32.6	35.1	40.9	41.0	42.6

Table 4: **DADA performance when trained on fractions of SYNTHIA.** Performance on Cityscapes as a function of the used percentage of training set.

remains loose as the adversarial losses are separately optimized. Our depth-aware adaptation framework S7 employing both feature fusion and DADA fusion performs best: paying more attention to closer objects during adversarial training is beneficial.

Annotation effort advantage. Table 4 reports DADA’s performance when trained on different fractions of the source dataset. Using only 50% of the SYNTHIA images with segmentation and depth annotations, DADA achieves comparable performance to AdvEnt trained on all images with segmentation annotations (40.9% vs. 40.8%). This finding is of practical importance for real-world setups where the source domain is also composed of real scenes: while dense depth annotation remains automatic in this case (through stereo matching as in Cityscapes or densification of sparse LiDAR measurements), semantic annotation must be manual, which incurs high costs and quality problems. Annotating fewer scenes can thus be beneficial even if depth is additionally required.

Limitations. We observe a few failure cases where different objects are indistinguishable due to blurry depth outputs. Improving depth quality may help in such cases. However, in our framework depth regression is only an auxiliary task which helps leveraging geometry-specific information to enrich visual representation and thus to improve on the main task. As in [24], paying too much attention to the auxiliary task actually hurts the performance on the main task.

5. Conclusion

In this work, we propose a novel UDA framework coined DADA – Depth-Aware Domain Adaptation – which leverages depth in source data as privileged information to help semantic segmentation on the target domain. This additional information is exploited through an auxiliary depth-prediction task that allows in turn a feature enrichment via fusion as well as a depth-aware modification of the original adaptation loss. Our experimental evaluations show that DADA consistently outperforms other UDA methods on different *synthetic-2-real* semantic segmentation benchmarks. As a direction of future work, we envisage the extension to real-world scenarios where depth information in the source domain is only sparsely available, *e.g.*, as provided by automotive laser scanners (LiDARs).

References

- [1] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT*. 2010. 5
- [2] Lin Chen, Wen Li, and Dong Xu. Recognizing RGB images by learning from RGB-D data. In *CVPR*, 2014. 3
- [3] 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. *PAMI*, 2018. 5
- [4] Yuhua Chen, Wen Li, and Luc Van Gool. Road: Reality oriented adaptation for semantic segmentation of urban scenes. In *CVPR*, 2018. 2
- [5] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *ICCV*, 2017. 2
- [6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 5
- [7] Gabriela Csurka. A comprehensive survey on domain adaptation for visual applications. In *Domain Adaptation in Computer Vision Applications*, pages 1–35. Springer, 2017. 1, 2
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5
- [9] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015. 2
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [12] Judy Hoffman, Saurabh Gupta, and Trevor Darrell. Learning with side information through modality hallucination. In *CVPR*, 2016. 3
- [13] 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 *ICML*, 2018. 2
- [14] 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
- [15] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, 2018. 3
- [16] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [17] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3DV*, 2016. 3
- [18] Kuan-Hui Lee, German Ros, Jie Li, and Adrien Gaidon. SPI-GAN: Privileged adversarial learning from simulation. In *ICLR*, 2019. 2, 3, 5, 6, 7
- [19] Wen Li, Li Niu, and Dong Xu. Exploiting privileged information from web data for image categorization. In *ECCV*, 2014. 3
- [20] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015. 2
- [21] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael Jordan. Unsupervised domain adaptation with residual transfer networks. In *NIPS*, 2016. 2
- [22] David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. Unifying distillation and privileged information. In *ICLR*, 2016. 3
- [23] 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. *arXiv preprint arXiv:1809.09478*, 2019. 2, 6
- [24] Taylor Mordan, Nicolas Thome, Gilles Henaff, and Matthieu Cord. Revisiting multi-task learning with rock: A deep residual auxiliary block for visual detection. In *NIPS*, 2018. 3, 8
- [25] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *CVPR*, 2018. 3
- [26] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kotschieder. The Mapillary Vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017. 5, 7
- [27] 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. In *NIPS Workshop*, 2017. 5
- [28] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016. 5
- [29] 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 *CVPR*, 2016. 2, 5
- [30] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. In *ICLR*, 2018. 2
- [31] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 2018. 2
- [32] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *CVPR*, 2018. 3
- [33] Nikolaos Sarafianos, Michalis Vrigkas, and Ioannis A Kakiaris. Adaptive SVM+: Learning with privileged information for domain adaptation. In *ICCV*, 2017. 3
- [34] Viktoriia Sharmanska, Novi Quadrianto, and Christoph Lampert. Learning to rank using privileged information. In *ICCV*, 2013. 3
- [35] 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 *CVPR*, 2018. 2, 4, 5, 6

- [36] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative representations. *arXiv preprint arXiv:1901.05427*, 2019. 2, 6
- [37] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017. 2
- [38] Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural networks*, 2009. 3
- [39] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. ADVENT: Adversarial entropy minimization for domain adaptation in semantic segmentation. *arXiv preprint arXiv:1811.12833*, 2019. 2, 4, 5, 6
- [40] Ziheng Wang and Qiang Ji. Classifier learning with hidden information. In *CVPR*, 2015. 3
- [41] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gokhan Uzunbas, Tom Goldstein, Ser Nam Lim, and Larry S. Davis. DCAN: Dual channel-wise alignment networks for unsupervised scene adaptation. In *ECCV*, 2018. 2
- [42] Hongliang Yan, Yukang Ding, Peihua Li, Qilong Wang, Yong Xu, and Wangmeng Zuo. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *CVPR*, 2017. 2
- [43] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *ICCV*, 2017. 3
- [44] Yiheng Zhang, Zhaofan Qiu, Ting Yao, Dong Liu, and Tao Mei. Fully convolutional adaptation networks for semantic segmentation. In *CVPR*, 2018. 3
- [45] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 2
- [46] Xinge Zhu, Hui Zhou, Ceyuan Yang, Jianping Shi, and Dahua Lin. Penalizing top performers: Conservative loss for semantic segmentation adaptation. In *ECCV*, 2018. 2
- [47] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018. 3