

Unsupervised Intra-domain Adaptation for Semantic Segmentation through Self-Supervision

Fei Pan Inkyu Shin Francois Rameau Seokju Lee In So Kweon
KAIST, South Korea

{feipan, dlsrbgg33, frameau, seokju91, iskweon77}@kaist.ac.kr

Abstract

Convolutional neural network-based approaches have achieved remarkable progress in semantic segmentation. However, these approaches heavily rely on annotated data which are labor intensive. To cope with this limitation, automatically annotated data generated from graphic engines are used to train segmentation models. However, the models trained from synthetic data are difficult to transfer to real images. To tackle this issue, previous works have considered directly adapting models from the source data to the unlabeled target data (to reduce the inter-domain gap). Nonetheless, these techniques do not consider the large distribution gap among the target data itself (intra-domain gap). In this work, we propose a two-step self-supervised domain adaptation approach to minimize the inter-domain and intra-domain gap together. First, we conduct the inter-domain adaptation of the model, from this adaptation, we separate target domain into an easy and hard split using an entropy-based ranking function. Finally, to decrease the intra-domain gap, we propose to employ a self-supervised adaptation technique from the easy to the hard subdomain. Experimental results on numerous benchmark datasets highlight the effectiveness of our method against existing state-of-the-art approaches. The source code is available at <https://github.com/feipan664/IntraDA.git>.

1. Introduction

Semantic segmentation aims at assigning each pixel in the image to a semantic class. Recently, convolutional neural network-based segmentation models [13, 33] have achieved remarkable progresses, leading to various applications in computer vision systems, such as autonomous driving [14, 31], robotics [15, 23], and disease diagnosis [35, 32]. Training such a segmentation network requires large amounts of annotated data. However, collecting large scale datasets with pixel-level annotations for semantic seg-




Figure 1: We propose a two-step self-supervised domain adaptation technique for semantic segmentation. Previous works solely **adapt** the segmentation model from the source domain to the target domain. Our work also consider **adapt-ing** from the clean map to the noisy map within the target domain.

mentation is difficult since they are expensive and labor intensive. Recently, photorealistic data rendered from simulators and game engines [19, 20] with precise pixel-level semantic annotations have been utilized to train segmentation networks. However, the models trained from synthetic data are hardly transferable to real data due to the cross-domain difference [11]. To address this issue, unsupervised domain adaptation (UDA) techniques have been proposed to align the distribution shift between the labeled source data and the unlabeled target data. For the particular task of semantic segmentation, adversarial learning-based UDA approaches demonstrate efficiency in aligning features at the image [16, 11] or output [26, 25] level.

More recently, the entropy of pixel-wise output predictions proposed by [28] is also used for output level alignment. Other approaches [38, 37] involve generating pseudo labels for target data and conducting refinement via an iterative self-training process. While many models consider the single-source-single-target adaptation setting, recent works [18, 34] have proposed to address the issue of multiple source domains; it focus on the multiple-source-single-target adaptation setting. Above all, previous works have mostly considered adapting models from the source data to the target data (inter-domain gap).

However, target data collected from the real world have diverse scene distributions; these distributions are caused by various factors such as moving objects, weather conditions, which leads to a large gap in the target (intra-domain gap). For example, the noisy map and clean map in the target domain, shown in Figure 1, are the predictions made by the same model on different images. While previous studies solely focus on reducing the inter-domain gap, the problem of the intra-domain gap has attracted a relatively low attention. In this paper, we present a two-step domain adaptation approach to minimize the inter-domain and intra-domain gap. Our model consists of three parts, which are presented in Figure 2, namely, 1) an inter-domain adaptation module to close inter-domain gap between the labeled source data and the unlabeled target data, 2) an entropy-based ranking system to separate target data into the an easy and hard split, and 3) an intra-domain adaptation module to close intra-domain gap between the easy and hard split (using pseudo labels from the easy subdomain). For semantic segmentation, our proposed approach achieves good performance against state-of-the-art approaches on benchmark datasets. Furthermore, our approach outperforms previous domain adaptation approaches for digit classification.

The Contributions of Our Work. First, we introduce the inter-domain gap among target data and propose an entropy-based ranking function to separate target domain into an easy and hard subdomain. Second, we propose a two-step self-supervised domain adaptation approach to minimize the inter-domain and intra-domain gap together.

2. Related Works

Unsupervised Domain Adaptation. The goal of unsupervised domain adaptation is to align the distribution shift between the labeled source and the unlabeled target data. Recently, adversarial-based UDA approaches have shown great capabilities in learning domain invariant features, even for complex tasks like semantic segmentation [28, 4, 26, 25, 21, 11, 17]. Adversarial-based UDA models for semantic segmentation usually involve two networks. One network is used as a generator to predict

the segmentation maps of input images, which can be from the source or the target. Given features from the generator, the second network functions as a discriminator to predict the domain labels. The generator tries to fool the discriminator, so as to align the distribution shift of the features from the two domains. Besides feature level alignment, other approaches try to align domain shift at the image level or output level. At the image level, CycleGAN [36] was applied in [11] to build generative images for domain alignment. At the output level, [25] proposes an end-to-end model involving structural output alignment for distribution shift. More recently, [28] takes advantage of the entropy of pixel-wise predictions from the segmentation outputs to address the domain gap. While all the previous studies exclusively considered aligning the inter-domain gap, our approach further minimizes the intra-domain gap. Thus, our technique can be combined with most existing UDA approaches for extra performance gains.

Uncertainty via Entropy. Uncertainty measurement has a strong connection with unsupervised domain adaptation. For instance, [28] proposes minimizing the target entropy value of the model outputs directly or using adversarial learning [25, 11] to close the domain gap for semantic segmentation. Also the entropy of the model outputs [29] is used as a confidence measurement for transferring samples across domains [24]. We propose utilizing entropy to rank target images to separate them into two an easy and hard split.

Curriculum Domain Adaptation. Our work is also related to curriculum domain adaptation [22, 31, 7] which deals with easy samples first. For curriculum domain adaptation on foggy scene understanding, [22] proposes to adapt a semantic segmentation model from non-foggy images to synthetic light foggy images, and then to real heavy foggy images. To generalize this concept, [7] decomposes the domain discrepancy into multiple smaller discrepancies by introducing unlabeled intermediate domains. However, these techniques require additional information to decompose domains. To cope with this limitation, [31] focuses on learning the global and local label distributions of images as the first task to regularize the model predictions in the target domain. In contrast, we propose a simpler and data-driven approach to learn the easy target samples based on an entropy ranking system.

3. Approach

Let \mathcal{S} denote a source domain consisting of a set of images $\subset \mathbb{R}^{H \times W \times 3}$ with their associated ground-truth C -class segmentation maps $\subset (1, C)^{H \times W}$; similarly, let \mathcal{T} denote a target domain containing a set of unlabeled images $\subset \mathbb{R}^{H \times W \times 3}$. In this section, a two-step self-supervised do-

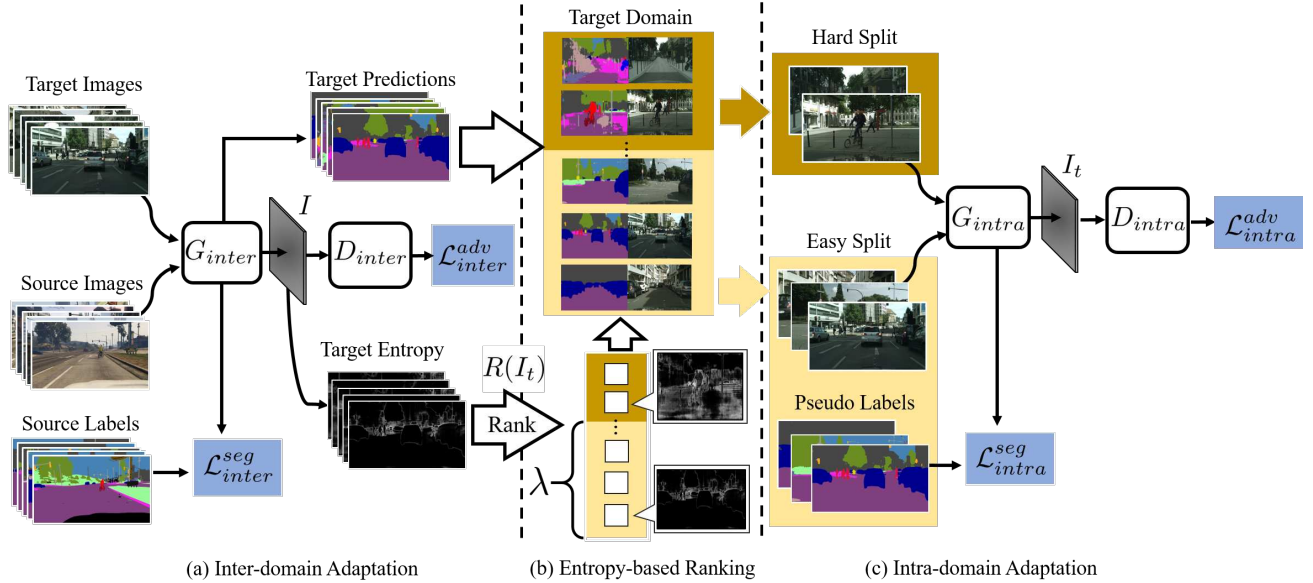


Figure 2: The proposed self-supervised domain adaptation model contains the inter-domain generator and discriminator $\{G_{inter}, D_{inter}\}$, and the intra-domain generator and discriminator $\{G_{intra}, D_{intra}\}$. The proposed model consists of three parts, namely, (a) an inter-domain adaptation, (b) an entropy-based ranking system, and (c) an intra-domain adaptation. In (a), given the source and the unlabeled target data, D_{inter} is trained to predict the domain label for the samples while G_{inter} is trained to fool D_{inter} . $\{G_{inter}, D_{inter}\}$ are optimized by minimizing the segmentation loss $\mathcal{L}_{inter}^{seg}$ and the adversarial loss $\mathcal{L}_{inter}^{adv}$. In (b), an entropy based function $R(I_t)$ is used to separate all target data into easy split and hard split. A hyperparameter λ is introduced as a ratio of target images assigned into the easy split. In (c), an intra-domain adaptation is used to close the gap between easy split and hard split. The segmentation predictions of easy split data from G_{inter} serve as pseudo labels. Given easy split data with pseudo labels and hard split data, D_{intra} is used to predict whether the sample is from easy split or hard split, while G_{intra} is trained to confuse D_{intra} . $\{G_{intra}$ and $D_{intra}\}$ are optimized using the intra-domain segmentation loss $\mathcal{L}_{intra}^{seg}$ and the adversarial loss $\mathcal{L}_{intra}^{adv}$.

main adaptation for semantic segmentation is introduced. The first step is the inter-domain adaptation, which is based on common UDA approaches [28, 25]. Then, the pseudo labels and predicted entropy maps of target data are generated such that the target data can be clustered into an easy and hard split. Specifically, an entropy-based ranking system is used to cluster the target data into the easy and hard split. The second step is the intra-domain adaptation, which consists in aligning the easy split with pseudo labels to the hard split, as shown in Figure 2. The proposed network consists of the inter-domain generator and discriminator $\{G_{inter}, D_{inter}\}$, and the intra-domain generator and discriminator $\{G_{intra}, D_{intra}\}$.

3.1. Inter-domain Adaptation

A sample $X_s \in \mathbb{R}^{H \times W \times 3}$ is from the source domain with its associated map Y_s . Each entry $Y_s^{(h,w)} = [Y_s^{(h,w,c)}]_c$ of Y_s provides a label of a pixel (h, w) as a one-hot vector. The network G_{inter} takes X_s as an input and generates a “soft-segmentation map” $P_s = G_{inter}(X_s)$.

Each C -dimensional vector $[P_s^{(h,w,c)}]_c$ at a pixel (h, w) serves as a discrete distribution over C classes. Given X_s with its ground-truth annotation Y_s , G_{inter} is optimized in a supervised way by minimizing the cross-entropy loss:

$$\mathcal{L}_{inter}^{seg}(X_s, Y_s) = - \sum_{h,w} \sum_c Y_s^{(h,w,c)} \log(P_s^{(h,w,c)}). \quad (1)$$

To close the inter-domain gap between the source and target domains, [28] proposes to utilize entropy maps in order to align the distribution shift of the features. The assumption of [28] is that the trained models tend to produce over-confident (low-entropy) predictions for source-like images, and under-confident (high-entropy) predictions for target-like images. Due to its simplicity and effectiveness, [28] is adopted in our work to conduct inter-domain adaptation. The generator G_{inter} takes a target image X_t as an input and produce the segmentation map $P_t = G_{inter}(X_t)$; the entropy map I_t is formulated as:

$$I_t^{(h,w)} = \sum_c -P_t^{(h,w,c)} \cdot \log(P_t^{(h,w,c)}). \quad (2)$$

To align the inter-domain gap, D_{inter} is trained to predict the domain labels for the entropy maps, while G_{inter} is trained to fool D_{inter} ; the optimization of G_{inter} and D_{inter} is achieved via the following loss function:

$$\mathcal{L}_{inter}^{adv}(X_s, X_t) = - \sum_{h,w} \log(1 - D_{inter}(I_s^{(h,w)})) + \log(D_{inter}(I_t^{(h,w)})), \quad (3)$$

where I_s is the entropy map of X_s . The loss function $\mathcal{L}_{inter}^{adv}$ and $\mathcal{L}_{inter}^{seg}$ are optimized to align the distribution shift between the source and target data. However, there remains a need for an efficient method that can minimize the intra-domain gap. For this purpose, we propose to separate the target domain into an easy and hard split and to conduct an intra-domain adaptation.

3.2. Entropy-based Ranking

Target images collected from the real world have diverse distributions due to various weather conditions, moving objects, and shading. In Figure 2, some target prediction maps are clean¹ and others are very noisy, despite being generated from the same model. Since the intra-domain gap exists among target images, a straightforward solution is to decompose the target domain into small subdomains/splits. However, it remains a challenging task due to the lack of target labels. To build these splits, we take advantage of the entropy maps in order to determine the confidence levels of the target predictions. The generator G_{inter} takes a target image X_t as input to generate P_t and the entropy map I_t . On this basis, we adopt a simple yet effective way for ranking by using:

$$R(X_t) = \frac{1}{HW} \sum_{h,w} \sum_c I_t^{(h,w,c)}, \quad (4)$$

which is the mean value of entropy map I_t . Given a ranking of scores from $R(X_t)$, hyperparameter λ is introduced as a ratio to separate the target images into an easy and hard split. Let X_{te} and X_{th} denote a target image assigned to the easy and hard split, respectively. In order to conduct domain separation, we define $\lambda = \frac{|X_{te}|}{|X_t|}$, where $|X_{te}|$ is the cardinality of the easy split, and $|X_t|$ is the cardinality of the whole target image set. To access the influence of λ , we conduct an ablation study on how to optimize λ in Table 3. Note that we do not introduce a hyperparameter as the threshold value for separation. The reason is that the threshold value is dependent on a specific dataset. However, we choose a hyperparameter as ratio, which shows strong generalization to other datasets.

¹The prediction map is clean means that the prediction is confident and smooth.

3.3. Intra-domain Adaptation

Since no annotation is available for the easy split, directly aligning the gap between the easy and hard split is infeasible. But we propose to utilize the predictions from G_{inter} as pseudo labels. Given an image from the easy split X_{te} , we forward X_{te} to G_{inter} and obtain the prediction map $P_{te} = G_{inter}(X_{te})$. While P_{te} is a ‘‘soft-segmentation map’’, we convert P_{te} to \mathcal{P}_{te} where each entry is a one-hot vector. With the aid of pseudo labels, G_{intra} is optimized by minimizing the cross-entropy loss:

$$\mathcal{L}_{intra}^{seg}(X_{te}) = - \sum_{h,w} \sum_c \mathcal{P}_{te}^{(h,w,c)} \log(G_{intra}(X_{te})^{(h,w,c)}). \quad (5)$$

To bridge the intra-domain gap between easy and hard split, we adopt the alignment on the entropy map for both splits. An image X_{th} from hard split is taken as input to the generator G to generate the segmentation map $P_{th} = G(X_{th})$ and the entropy map I_{th} . To close the inter-domain gap, the intra-domain discriminator D_{intra} is trained to predict the split labels of I_{te} and I_{th} : I_{te} is from the easy split, and I_{th} is from the hard split. G is trained to fool D_{intra} . The adversarial learning loss to optimize G_{intra} and D_{intra} is formulated as:

$$\mathcal{L}_{intra}^{adv}(X_{te}, X_{th}) = - \sum_{h,w} \log(1 - D_{intra}(I_{te}^{(h,w)})) + \log(D_{intra}(I_{th}^{(h,w)})). \quad (6)$$

Finally, our complete loss function \mathcal{L} is formed by all the loss functions:

$$\mathcal{L} = \mathcal{L}_{inter}^{seg} + \mathcal{L}_{inter}^{adv} + \mathcal{L}_{intra}^{seg} + \mathcal{L}_{intra}^{adv}, \quad (7)$$

and our objective is to learn a target model G according to:

$$G^* = \arg \min_{G_{intra}} \min_{G_{inter}} \max_{D_{inter}} \mathcal{L}. \quad (8)$$

Since our proposed model is two-step self-supervised approach, it is difficult to minimize \mathcal{L} in one training stage. Thus, we choose to minimize it in three stages. First, we train the inter-domain adaptation for the model to optimize G_{inter} and D_{inter} . Second, we generate target pseudo labels by utilizing G_{inter} and rank all target images based on $S(X_t)$. Finally, we train the intra-domain adaptation to optimize G_{intra} and D_{intra} .

4. Experiments

In this section, we introduce the experimental details of the inter-domain and the intra-domain adaptation on semantic segmentation.

Table 1: The semantic segmentation results of Cityscapes validation set with models trained on GTA5 (a), SYNTHIA (b), and Synscapes (c). All the results are generated from the ResNet-101-based models. In the experiments of (a) and (b), AdvEnt[28] is used as the framework for the inter-domain adaptation and intra-domain adaptation. In the experiment of (c), AdaptSegNet [25] is used as the framework of the inter-domain adaptation and intra-domain adaptation. mIoU* in (b) denotes the mean IoU of 13 classes, excluding the classes with *.

(a) GTA5 → Cityscapes																				
Method	road	sidewalk	building	wall	fence	pole	light	sign	veg	terrain	sky	person	rider	car	truck	bus	train	mbike	bike	mIoU
Without adaptation [25]	75.8	16.8	77.2	12.5	21.0	25.5	30.1	20.1	81.3	24.6	70.3	53.8	26.4	49.9	17.2	25.9	6.5	25.3	36.0	36.6
ROAD [5]	76.3	36.1	69.6	28.6	22.4	28.6	29.3	14.8	82.3	35.3	72.9	54.4	17.8	78.9	27.7	30.3	4.0	24.9	12.6	39.4
AdaptSegNet [25]	86.5	36.0	79.9	23.4	23.3	23.9	35.2	14.8	83.4	33.3	75.6	58.5	27.6	73.7	32.5	35.4	3.9	30.1	28.1	42.4
MinEnt [28]	84.2	25.2	77.0	17.0	23.3	24.2	33.3	26.4	80.7	32.1	78.7	57.5	30.0	77.0	37.9	44.3	1.8	31.4	36.9	43.1
AdvEnt [28]	89.9	36.5	81.6	29.2	25.2	28.5	32.3	22.4	83.9	34.0	77.1	57.4	27.9	83.7	29.4	39.1	1.5	28.4	23.3	43.8
Ours	90.6	37.1	82.6	30.1	19.1	29.5	32.4	20.6	85.7	40.5	79.7	58.7	31.1	86.3	31.5	48.3	0.0	30.2	35.8	46.3

(b) SYNTHIA → Cityscapes																				
Method	road	sidewalk	building	wall*	fence*	pole*	light	sign	veg	sky	person	rider	car	bus	mbike	bike	mIoU	mIoU*		
Without adaptation [25]	55.6	23.8	74.6	9.2	0.2	24.4	6.1	12.1	74.8	79.0	55.3	19.1	39.6	23.3	13.7	25.0	33.5	38.6		
AdaptSegNet [25]	81.7	39.1	78.4	11.1	0.3	25.8	6.8	9.0	79.1	80.8	54.8	21.0	66.8	34.7	13.8	29.9	39.6	45.8		
MinEnt [28]	73.5	29.2	77.1	7.7	0.2	27.0	7.1	11.4	76.7	82.1	57.2	21.3	69.4	29.2	12.9	27.9	38.1	44.2		
AdvEnt [28]	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		
Ours	84.3	37.7	79.5	5.3	0.4	24.9	9.2	8.4	80.0	84.1	57.2	23.0	78.0	38.1	20.3	36.5	41.7	48.9		

(c) Synscapes → Cityscapes																				
Method	road	sidewalk	building	wall	fence	pole	light	sign	veg	terrain	sky	person	rider	car	truck	bus	train	mbike	bike	mIoU
Without adaptation	81.8	40.6	76.1	23.3	16.8	36.9	36.8	40.1	83.0	34.8	84.9	59.9	37.7	78.4	20.4	20.5	7.8	27.3	52.5	45.3
AdaptSegNet [25]	94.2	60.9	85.1	29.1	25.2	38.6	43.9	40.8	85.2	29.7	88.2	64.4	40.6	85.8	31.5	43.0	28.3	30.5	56.7	52.7
Ours	94.0	60.0	84.9	29.5	26.2	38.5	41.6	43.7	85.3	31.7	88.2	66.3	44.7	85.7	30.7	53.0	29.5	36.5	60.2	54.2

4.1. Datasets

In the experiments of semantic segmentation, we adopt the setting of adaptation from the synthetic to the real domain. To conduct this series of tests, synthetic datasets including GTA5 [19], SYNTHIA [20] and Synscapes [30] are used as source domains, along with the real-world dataset Cityscapes [6] as the target domain. Models are trained given labeled source data and unlabeled target data. Our model is evaluated on Cityscapes validation set.

- **GTA5:** The synthetic dataset GTA5 [19] contains 24,966 synthetic images with a resolution of 1914 × 1052px and corresponding ground-truth annotations. These synthetic images are collected from a video game based on the urban scenery of Los Angeles city. The ground-truth annotations generated automatically contain 33 categories. For training, we consider only 19 categories which are compatible with the Cityscapes dataset [6], similarly to previous work.
- **SYNTHIA:** SYNTHIA-RAND-CITYSCAPES [20] is used as another synthetic dataset. It contains 9,400 fully annotated RGB images. During the training time, we consider the 16 common categories with the

Cityscapes dataset. During evaluation, 16- and 13-class subsets are used to evaluate the performance.

- **Synscapes:** Synscapes [30] is a photorealistic synthetic dataset consisting of 25,000 fully annotated RGB images with a resolution of 1440 × 720px. Alike Cityscape, the ground-truth annotations contain 19 categories.
- **Cityscapes:** As the dataset collected from real world, Cityscapes [6] provides 3975 images with fine segmentation annotations. 2975 images are taken from the training set of Cityscapes to be used for training. The 500 images from the evaluation set of Cityscapes are used to evaluate the performance of our model.

Evaluation. The semantic segmentation performance is evaluated on every category using the PASCAL VOC intersection-over-union metric, i.e., $\text{IoU} = \text{TP}/(\text{TP} + \text{FP} + \text{FN})$ [9], where TP, FP, and FN are the number of true positive, false positive, and false negative pixels, respectively.

Implementation Details. In the experiments for GTA5→Cityscapes and SYNTHIA→Cityscapes, we utilize

Table 2: The self-training and intra-domain adaptation gain on GTA5 \rightarrow Cityscapes.

Model	mIoU
AdvEnt [25]	43.8
AdvEnt + intra-domain adaptation	45.1
AdvEnt self-training ($\lambda = 1.0$)	45.5
Ours	46.3
Ours + entropy normalization	47.0

Table 3: The ablation study on hyperparameter λ for separating the target domain into the easy and the hard split.

GTA5 \rightarrow Cityscapes						
λ	0.0	0.5	0.6	0.67	0.7	1.0
mIoU	43.8	45.2	46.0	46.3	45.6	45.5

the framework of AdvEnt [28] to train G_{inter} and D_{inter} for inter-domain adaptation; the backbone of G_{inter} is a ResNet-101 architecture [10] with pretrained parameters from ImageNet [8]; the input data are labeled source images and unlabeled target images. The model for inter-domain adaptation G_{inter} is trained for 70,000 iterations. After training, G_{inter} is used to generate the segmentation and entropy maps for all 2,975 images from Cityscapes training set. Then, we utilize $R(I_t)$ to get the ranking scores for all target images and to separate them into the easy and hard split based on λ . We conduct an ablation study of λ for optimization in Table 3. For the intra-domain adaptation, G_{inter} has same architecture as G_{inter} , and D_{intra} same as D_{inter} ; the input data are 2,975 Cityscapes training images with pseudo labels of easy split. G_{intra} is trained with the pretrained parameters from ImageNet and D_{intra} from scratch, similar to AdvEnt. In addition to the previously mentioned experiments, we also conduct the experiment for Synscapes \rightarrow Cityscapes. For comparison with AdaptSegNet [25], We apply the framework of AdaptSegNet in the experiment for the inter-domain and intra-domain adaptation.

Similarly to [28] and [25], we utilize the multi-level feature outputs from *conv4* and *conv5* for inter-domain adaptation and intra-domain adaptation. To train G_{inter} and G_{intra} , we apply an SGD optimizer [2] with a learning rate of 2.5×10^{-4} , momentum 0.9, and a weight decay 10^{-4} for training G_{inter} and G_{intra} . An Adam optimizer [12] with a learning rate of 10^{-4} is used for training D_{inter} and D_{intra} .

4.2. Results

GTA5. In Table 1 (a), we compare the segmentation performance of our method with other state-of-the-art methods [25, 5, 28] on Cityscapes validation set. For a fair comparison, the baseline model is adopted from DeepLab-

v2[3] with a ResNet-101 backbone. Overall, our proposed method achieves 46.3% in mean IoU. Compared to AdvEnt, the intra-domain adaptation from our method leads to a 2.5% improvement in mean IoU.

To highlight the relevance of the proposed intra-domain adaptation, we conduct a comparison with segmentation loss $\mathcal{L}_{intra}^{seg}$ and adversarial adaptation loss $\mathcal{L}_{intra}^{adv}$ in Table 2. The baseline AdvEnt [28] achieves 43.8% of mIoU. By using AdvEnt + intra-domain adaptation, which means $\mathcal{L}_{intra}^{seg} = 0$, we obtain 45.1%, showing the effectiveness of adversarial learning for the intra-domain alignment. By applying AdvEnt + self-training, with $\lambda = 1.0$ (all pseudo labels used for self-training), which means $\mathcal{L}_{intra}^{adv} = 0$, we achieve 45.5% of mIoU, underlying the importance of employing pseudo labels. Finally, our proposed model achieves 46.3% of mIoU (self-training + intra-domain alignment).

Admittedly, complex scenes (containing many objects) might be categorized as "hard". To provide a more representative "ranking", we adopt a new normalization by dividing the mean entropy with the number of predicted rare classes in the target image. For Cityscapes dataset, we define these rare classes as "wall, fence, pole, traffic light, traffic sign, terrain, rider, truck, bus, train, motor". The entropy normalization helps to move images with many objects to the easy split. By using the normalization, our proposed model achieves 47.0% of mIoU, as shown in Table 2. Our proposed method also has limitation for some classes.

In Figure 3, we provide some visualizations of segmentation maps from our technique. The segmentation maps generated from our model trained with inter-domain alignment and intra-domain alignment are more accurate than the baseline model AdvEnt, which has been only trained with inter-domain alignment. A representative set of images belonging to the "hard" split are visible in Figure 4. After intra-domain alignment, we produce the segmentation maps shown in the (d) column. Compared with (c) column, our model can be transferred to more difficult target images.

Analysis of Hyperparameter λ . We conduct a study on finding a proper value for the hyperparameter λ in our experiment of GTA5 \rightarrow Cityscapes. In Table 3, different values of δ are used for setting up the decision boundary for domain separation. When $\lambda = 0.67$, i.e., the ratio of $|X_{te}|$ to $|X_t|$ is approximately 2/3, the model achieves 46.3 of mIoU as the best performance on Cityscapes validation set.

SYNTIA. We use SYNTIA as the source domain and present evaluation results of the proposed method and state-of-the-art methods [25, 28] on Cityscapes validation set in the Table 1. For a fair comparison, we also adopt the same DeepLab-v2 with the ResNet-101 architecture. Our method is evaluated on both 16-class and 13-class

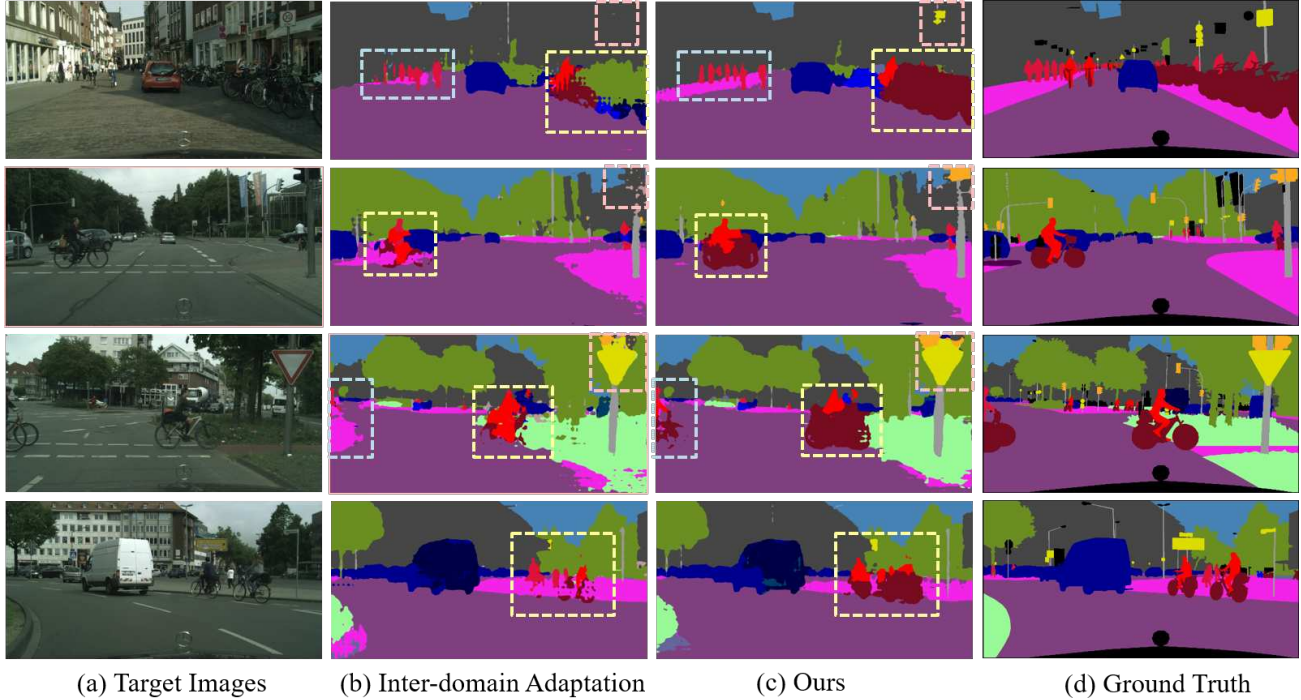


Figure 3: The example results of evaluation for GTA5→Cityscapes. (a) and (d) are the images from Cityscapes validation set and the corresponding ground-truth annotation. (b) are the predicted segmentation maps of the inter-domain adaptation [28]. (c) are the predicted maps from our technique.

baselines. According to the results in Table 1 (b), our proposed method has achieved 41.7% and 48.9% of mean IoU on 16-class and 13-class baseline, respectively. As shown in the Table 1, our model is significantly more accurate on the car and motor bike classes than existing techniques. The reason is that we apply the intra-domain adaptation to further narrow the domain gap.

Synscapes. The only work that we currently have found using Synscapes dataset is [25]. Thus we use AdaptSegNet [25] as our baseline model. In order to present a fair comparison, We only consider using vanilla-GAN in our experiments. With inter-domain and intra-domain adaptation, our model achieves 54.2% of mIoU, which higher than AdaptSegNet shown in Table 1 (c).

4.3. Discussion

Theoretical Analysis. Comparing (a), (b) in Table 1, GTA5 to Cityscapes is more effective than SYNTHIA to Cityscapes. We believe the reason is that GTA5 has more similar images of street scenes with Cityscapes than other synthetic datasets. We also provide a theoretical analysis. Let \mathcal{H} denote the hypothesis class, S and T be the source and the target domain. The theory from [1] proposes to bound the expected error on the target

domain $\epsilon_T(h)$: $\forall h \in \mathcal{H}, \epsilon_T(h) \leq \epsilon_S(h) + \frac{1}{2}d_{\mathcal{H}}(S, T) + \Lambda$, where $\epsilon_S(h)$ is the expected error on the source domain, $d_{\mathcal{H}}(S, T) = 2 \sup |\Pr_S(h) - \Pr_T(h)|$ is the distance for domain divergence, and Λ is considered as a constant in normal cases. Therefore, $\epsilon_T(h)$ is upper bounded by $\epsilon_S(h)$ and $d_{\mathcal{H}}(S, T)$ in our case. Our proposed model is to minimize $d_{\mathcal{H}}(S, T)$ by using the inter-domain and the intra-domain alignment together. If $d_{\mathcal{H}}(S, T)$ has high value, the higher upper bound in the first stage of the inter-domain adaptation affects our entropy ranking system, and the intra-domain adaptation processes. Therefore, our model is less efficient in big domain gap. With respect to the limitation, our model performance is affected by $d_{\mathcal{H}}(S, T)$ and $\epsilon_S(h)$. Firstly, the larger divergence of the source and target domain leads to higher value in $d_{\mathcal{H}}(S, T)$. The upper bound of error is higher so our model would be less effective. Secondly, $\epsilon_S(h)$ would be very high when the model uses small neural networks. In such case, our model would also be less effective.

Digit Classification. Our model are also capable to be applied in digit classification task. We consider the adaptation shift of MNIST→USPS, USPS→MNIST, and SVHN→MNIST. Our model is trained using the training sets: MNIST with 60000 images, USPS with 7291 images,




Figure 4: The examples of entropy maps from hard split for GTA5→Cityscapes. (a) are the hard images from Cityscapes training set. (b) and (c) are the predicted entropy and segmentation maps from model trained solely by the inter-domain adaptation [28]. (d) are the improved predicted segmentation results of the hard images from our model.

Table 4: The experimental results of adaptation across digit datasets.

Model	MNIST → USPS	USPS → MNIST	SVHN → MNIST
Source only	82.2±0.8	69.6±3.8	67.1±0.6
ADDA [27]	89.4±0.2	90.1±0.8	76.0±1.8
CyCADA [11]	95.6±0.2	96.5±0.1	90.4±0.4
Ours	95.8±0.1	97.8±0.1	95.1±0.3

standard SVHN with 73257 images. The proposed model is evaluated on the standard test sets: MNIST with 10000 images and USPS with 2007 images. In digit classification task, G_{inter} and G_{intra} serve as classifiers with same architecture, which is based on a variant of the LeNet architecture. In inter-domain adaptation, We utilize the framework of CyCADA [11] to train G_{inter} and D_{inter} . In the ranking stage, we utilize G_{inter} to generate the predictions of all target data and compute their ranking score using $R(X_t)$. With respect to λ , we adopted $\lambda = 0.8$ in all experiments. Our network for intra-domain adaptation is also based on CyCADA [11]. In Table 4, our proposed model achieve 95.8±0.1% of accuracy on MNIST → USPS, 97.8±0.1% of accuracy on USPS → MNIST, and 95.1±0.3% on SVHN → MNIST. Our model outperforms the baseline model CyCADA [11].

5. Conclusion

In this paper, we present a self-supervised domain adaptation to minimize the inter-domain and intra-domain gap simultaneously. We first train the model using the inter-domain adaptation from existing approaches. Secondly, we produce target image entropy maps and use an entropy-based ranking functions to split the target domain. Lastly, we conduct the intra-domain adaptation to further narrow the domain gap. We conduct extensive experiments on synthetic to real images in traffic scenarios. Our model can be combined with existing domain adaptation approaches. Experimental results shows that our model outperforms existing adaptation algorithms.

Acknowledgments

This research was partially supported by the Shared Sensing for Cooperative Cars Project funded by Bosch (China) Investment Ltd. This work was also partially supported by the Korea Research Fellowship Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT and Future Planning (2015H1D3A1066564).

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *NeurIPS*, pages 137–144, 2007. 7
- [2] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010. 6
- [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*, 40(4):834–848, 2017. 6
- [4] Minghao Chen, Hongyang Xue, and Deng Cai. Domain adaptation for semantic segmentation with maximum squares loss. In *ICCV*, pages 2090–2099, 2019. 2
- [5] Yuhua Chen, Wen Li, and Luc Van Gool. Road: Reality oriented adaptation for semantic segmentation of urban scenes. In *CVPR*, pages 7892–7901, 2018. 5, 6
- [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] Shuyang Dai, Kihyuk Sohn, Yi-Hsuan Tsai, Lawrence Carin, and Manmohan Chandraker. Adaptation across extreme variations using unlabeled domain bridges. *arXiv preprint arXiv:1906.02238*, 2019. 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*, pages 248–255. Ieee, 2009. 6
- [9] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *IJCV*, 111(1):98–136, 2015. 5
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [11] 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*, pages 1989–1998, 2018. 1, 2, 8
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [13] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 1
- [14] Pauline Luc, Natalia Neverova, Camille Couprie, Jakob Verbeek, and Yann LeCun. Predicting deeper into the future of semantic segmentation. In *ICCV*, pages 648–657, 2017. 1
- [15] Andres Milioto, Philipp Lottes, and Cyrill Stachniss. Real-time semantic segmentation of crop and weed for precision agriculture robots leveraging background knowledge in cnns. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2229–2235. IEEE, 2018. 1
- [16] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *CVPR*, pages 4500–4509, 2018. 1
- [17] Kwanyong Park, Sanghyun Woo, Dahun Kim, Donghyeon Cho, and In So Kweon. Preserving semantic and temporal consistency for unpaired video-to-video translation. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 1248–1257, 2019. 2
- [18] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, pages 1406–1415, 2019. 2
- [19] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *ECCV*, volume 9906 of *LNCS*, pages 102–118. Springer International Publishing, 2016. 1, 5
- [20] 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*, pages 3234–3243, 2016. 1, 5
- [21] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, pages 3723–3732, 2018. 2
- [22] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *ECCV*, pages 687–704, 2018. 2
- [23] Alexey A Shvets, Alexander Rakhlin, Alexandr A Kalinin, and Vladimir I Iglovikov. Automatic instrument segmentation in robot-assisted surgery using deep learning. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 624–628. IEEE, 2018. 1
- [24] Jong-Chyi Su, Yi-Hsuan Tsai, Kihyuk Sohn, Buyu Liu, Subhransu Maji, and Manmohan Chandraker. Active adversarial domain adaptation. In *WACV*, pages 739–748, 2020. 2
- [25] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, pages 7472–7481, 2018. 1, 2, 3, 5, 6, 7
- [26] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *ICCV*, pages 1456–1465, 2019. 1, 2
- [27] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017. 8
- [28] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *CVPR*, pages 2517–2526, 2019. 2, 3, 5, 6, 7, 8
- [29] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. Cost-effective active learning for deep image classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(12):2591–2600, 2016. 2

- [30] Magnus Wrenninge and Jonas Unger. Synscapes: A photorealistic synthetic dataset for street scene parsing. *arXiv preprint arXiv:1810.08705*, 2018. [5](#)
- [31] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *ICCV*, volume 2, page 6, Oct 2017. [1](#), [2](#)
- [32] Amy Zhao, Guha Balakrishnan, Fredo Durand, John V Guttag, and Adrian V Dalca. Data augmentation using learned transformations for one-shot medical image segmentation. In *CVPR*, pages 8543–8553, 2019. [1](#)
- [33] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. [1](#)
- [34] Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. In *NIPS*, pages 8559–8570, 2018. [2](#)
- [35] Yi Zhou, Xiaodong He, Lei Huang, Li Liu, Fan Zhu, Shanshan Cui, and Ling Shao. Collaborative learning of semi-supervised segmentation and classification for medical images. In *CVPR*, pages 2079–2088, 2019. [1](#)
- [36] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. [2](#)
- [37] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, pages 289–305, 2018. [2](#)
- [38] Yang Zou, Zhiding Yu, Xiaofeng Liu, B.V.K. Vijaya Kumar, and Jinsong Wang. Confidence regularized self-training. In *ICCV*, October 2019. [2](#)