

This WACV 2023 Workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Exploiting Inter-pixel Correlations in Unsupervised Domain Adaptation for Semantic Segmentation

Inseop Chung Jayeon Yoo Nojun Kwak Seoul National University

{jis3613, jayeon.yoo, nojunk}@snu.ac.kr

Abstract

"Self-training" has become a dominant method for semantic segmentation via unsupervised domain adaptation (UDA). It creates a set of pseudo labels for the target domain to give explicit supervision. However, the pseudo labels are noisy, sparse and do not provide any information about inter-pixel correlations. We regard inter-pixel correlation quite important because semantic segmentation is a task of predicting highly structured pixel-level outputs. Therefore, in this paper, we propose a method of transferring the inter-pixel correlations from the source domain to the target domain via a self-attention module. The module takes the prediction of the segmentation network as an input and creates a self-attended prediction that correlates similar pixels. The module is trained only on the source domain to learn the domain-invariant inter-pixel correlations, then later, it is used to train the segmentation network on the target domain. The network learns not only from the pseudo labels but also by following the output of the self-attention module which provides additional knowledge about the inter-pixel correlations. Through extensive experiments, we show that our method significantly improves the performance on two standard UDA benchmarks and also can be combined with recent state-of-the-art method to achieve better performance.

1. Introduction

Unsupervised domain adaptation (UDA) has become one of the primary methods to tackle label-scarce real-world problems such as simulated training for robots and autonomous driving. It transfers the knowledge of a deep neural network learned from the source domain with rich labels to the target domain without any label information. Usually, the source domain is photo-realistic synthetic dataset [19, 20] and the target domain is real world images [6]. UDA is especially useful for semantic segmentation since it requires dense pixel-level labeling process which is laborious and time-consuming. Recently, the "self-training" scheme has become a dominant method to enable UDA for semantic segmentation. This paper also tackles the UDA for semantic segmentation based on the "self-training" scheme, yet we introduce a novel method for transferring inter-pixel correlations in the source domain to the target domain, which is also simple and effective.

The idea of self-training [43, 44, 15] is to give supervision for the unlabeled target domain as well. It generates a set of pseudo labels for the target domain using a trained model and then retrains a new model with the generated pseudo labels. To enhance the performance, it iteratively repeats this process with the retrained model. However, since pseudo labels are based on confident pixels whose predicted confidence surpasses a certain threshold, they are not always correct and sometimes very sparse which means that some pixels have no pseudo labels. To resolve this issue, recent works [21, 36, 40] have proposed to generate better pseudo labels. They denoise, rectify, and densify the pseudo labels to give more correct supervision. However, these methods are somewhat heuristic and require complex algorithms. Also, pseudo labels provide information only about which class a pixel should be predicted as but not about how each pixel should be correlated or similar to other pixels. We argue that providing guidance about the correlation between pixels is important since the outputs of a semantic segmentation task are highly structured and correlated.

A semantic segmentation task requires a model to predict highly structured pixel-level outputs which contains densely related spatial and local information. For example, the pixels corresponding to the *sky* are usually located at an upper area of an image, *car* pixels are over the *road* pixels and *sidewalks* and *fence* pixels are usually on the sides of the *road*. These inter-pixel relations are domain-invariant properties which are consistent for different domains. Therefore, in this paper, we propose an idea of transferring this domain-invariant knowledge of the inter-pixel correlation from the source domain to the target domain. Our intuition is that providing the extra knowledge of domain-invariant pixel correlations can help to overcome and complement the defects of sparse and incorrect pseudo labels.

To effectively capture these inter-pixel correlations and transfer them to the target domain, we propose a novel selfattention [4, 18, 24] module. The proposed self-attention module takes predictions of the segmentation network and outputs new self-attended predictions. The module creates an attention map which captures similarities between pixels and uses this map to calculate a weighted sum of the given predictions. Our proposed algorithm consists of two steps. The first step is to train the self-attention module only on the source domain since it has the ground truth labels that can correctly guide about the domain-invariant inter-pixel correlations. The second step is to train the actual segmentation network using the trained self-attention module by transferring its learned knowledge to the target domain. At this phase, the attention module is frozen and the main segmentation network is the only network being trained. The module takes the predictions of the segmentation network as inputs and outputs new self-attended predictions. Then, we train the segmentation network to follow the new selfattended predictions. We call this the 'self-attention loss'. Overall, the segmentation network is trained with both the self-attention loss and the self-training loss. Detailed training methodology will be explained in Sec. 4.

We conduct extensive experiments on the two standard UDA benchmarks, GTA5-Cityscapes and SYN-THIA \rightarrow Cityscapes. We show that the network trained with our method outperforms the one trained only with selftraining. Especially, our method shows significant improvements on rare classes which scarcely appear in the datasets. Also, our method does not require a complex humandesigned algorithm, it rather lets the module to decide which knowledge to be learned and transferred, thus our method is simple and needs less human intervention. Moreover, our method can be applied to other state-of-the-art self-training UDA methods and additionally boost the performance. We combined our method with ProDA [36], one of the recent SOTA methods, and could observe a performance improvement which sets new SOTA scores on both UDA benchmarks tested.

2. Related Work

2.1. Unsupervised Domain Adaptation

Unsupervised domain adaptation has been widely studied to address the domain gap between labeled source data and unlabeled target data. Since domain adaptation is an important problem in real world environments, it has been studied in various tasks such as classification[8] and segmentation[43, 44, 36]. There are three major lines of domain adaptation approaches: adversarial learning, image translation, and self-training. Based on the theoretical analysis[8], the works [23, 26] adapt the segmentation model to the target domain via adversarial learning which tries to increase the domain discrimination loss of the discriminator while reducing the segmentation loss on the source domain. The works [11, 35] translate the labeled source domain images into the target domain styles using CycleGAN [42] and Fourier transforms to encourage image-level adaptation.

2.2. Self-training

Another line of unsupervised domain adaptation is selftraining which is first introduced in semantic segmentation by [43]. In [43], the authors generate pseudo labels using a model trained on the source domain and retrain the model with the pseudo labeled target domain images to implicitly encourage class-wise feature alignment and enable adaptation to the target domain. However, since pseudo labels are noisy due to domain discrepancy, many methods have been proposed to address this problem. [44] introduces a regularization method to generate soft pseudo labels and smooth the output of the model to avoid overfitting on overconfident labels. [36] and [40] lower the weights for noisy pseudo labels based on the feature distance from the class prototypes and the variance between the outputs of the two segmentation models respectively. In addition, [9] has the same effect of modifying pseudo labels based on prior distribution by introducing a transition probability map between the ground-truth label and the pseudo label. To address the sparsity of confident pseudo labels which is another limitation of self-training, [21] proposes densification of pseudo labels using confident pseudo labels of neighboring pixels. In this paper, we resolve these issues by transferring inter-pixel correlation which is a domain-invariant property.

2.3. Self-attention for vision

Self-attention calculates the response at a position in a sequence (*e.g.*, an image) by interacting with all positions and paying more attention to more salient positions. There have been several attempts to introduce attention modules into convolutional networks. Non-local network [29] captures long-range dependencies by computing interactions between any two positions regardless of their positional distance. [32] refines features by using channel and spatial attention for various vision tasks. [30] utilizes self-attention mechanism to refine CAM [41] for weakly supervised semantic segmentation. Recently, transformer[25] based models, which consist of attention mechanisms rather than convolutional architectures, are also being actively studied in vision tasks [2, 7].

3. Preliminary

In the common UDA setting, it is assumed that the dataset in the source domain has labels, $\{x_s, y_s\}$, while that



Figure 1. The architecture of the proposed self-attention module. Grey shaded area at the center is the self-attention module and \bigotimes refers to a matrix product. The module takes the output of the segmentation network and creates an attention map to apply self-attention to *z*. *z'* is the new self-attended output of the module which possesses information about inter-pixel correlations.

in the target domain is unlabelled, $\{x_t\}$. The main purpose of UDA for semantic segmentation is to train a segmentation network, \mathcal{G} , to perform well on the target domain by transferring the knowledge it has learned from the source domain. The two domains are assumed to share the same C classes, hence it is possible to adapt between them. Typically, for the source domain, the categorical cross-entropy loss is used to train the network utilizing the given ground truth labels:

$$\mathcal{L}_{seg}^{S}(x_{s}) = -\sum_{\substack{i=1\\(i,c)}}^{H \times W} \sum_{c=1}^{C} y_{s}^{(i,c)} \log(p_{s}^{(i,c)}).$$
(1)

Here, $p_s = \mathcal{G}(x_s)$ and $p_s^{(i,c)}$ refers to the softmax probability $p_s^{(i)}$ of pixel *i* belonging to class *c*. The label $y_s^{(i)}$ is a one-hot vector where $y_s^{(i,c)}$ is 1 if the ground truth class at pixel *i* is *c*, otherwise it is 0. However, for the target domain, the labels do not exist, thus pseudo labels $\{\hat{y}_t\}$ are generated to perform self-training. The pseudo labels are based on the confident pixels predicted by a trained network. Following the process of [15], the pseudo labels are assigned only to pixels with confidence higher than a given threshold:

$$\hat{y}_{t}^{(i,c)} = \begin{cases} 1, & \text{if } c = \arg\max_{c'} p_{t}^{(i,c')} \\ & \text{and } \max_{c'} p_{t}^{(i,c')} > \tau^{c} \\ 0, & \text{otherwise} \end{cases}$$
(2)

$$\mathcal{L}_{seg}^{T}(x_{t}) = -\sum_{i=1}^{H \times W} \sum_{c=1}^{C} \hat{y}_{t}^{(i,c)} \log(p_{t}^{(i,c)}).$$
(3)

The threshold τ^c is set differently for each class c and how it is determined is detailed in the supplementary. Using the generated pseudo labels, the target domain can also be trained in a supervised fashion as well using (3). The overall segmentation loss is the sum of (1) and (3). In this way, the performance of the network on the target domain can be greatly improved, however, in this work, we further enhance this by introducing our novel self-attention loss term.

4. Method

4.1. Self-attention Module (SAM)

A segmentation network uses an upsampling to enlarge its output to the size of its input. DeepLabV2 [3] uses bilinear interpolation while FCN-8s [16] uses transposed convolution. Therefore, it can be formulated as p = S(U(z)) =

 $\mathcal{G}(x) \in \mathbb{R}^{H \times W \times C}$ where S and U stands for the softmax and the upsampling respectively. The $z \in \mathbb{R}^{h imes w imes C}$ is the output of the network right before the upsampling (h, w <H, W). It is the logit which becomes the class probability distribution after softmax is applied. As shown in Fig. 1, z is inputted to our self-attention module to create a new selfattended output z'. We assume that z is flattened in the spatial dimension for simplicity, so its shape is $\mathbb{R}^{hw \times C}$. Then, z consists of hw number of logits, $z = (z_1, z_2, z_3, \dots z_{hw})^{\mathsf{T}}$ where each logit, z_i , is a C-dimensional vector. The architecture of the self-attention module is depicted in Fig. 1. First, z is transformed into another space using a 1×1 convolutional layer producing $\text{Conv}(z) \in \mathbb{R}^{hw \times \overline{C}}$. Then we calculate the attention map by outer-producting Conv(z) to itself. More specifically, if we denote Conv(z) as A to keep the notation uncluttered, the attention map is computed as

$$M = \frac{A \cdot A^{\mathsf{I}}}{\|A\|_2 \cdot \|A\|_2^{\mathsf{T}}} \in \mathbb{R}^{hw \times hw},\tag{4}$$

where \cdot is a matrix product, $||A||_2 \in \mathbb{R}^{hw \times 1}$ refers to the row-wise L2-norm of A along with its channel dimension and the division is done element-wise. Each row of M shows how each pixel is similar to other pixels in cosine similarity, hence its element is normalized between -1 and 1. We apply ReLU [1] activation on M to eliminate any negative similarity and then L1 normalize each row in order to make the sum of all elements in the same row equals 1. We denote M' as the attention map after the ReLU activation and the normalization. $M'_{[i,j]}$ refers to an element of *i*th row and *j*th column of M' and is computed as follows:

$$M'_{[i,j]} = \frac{ReLU(M)_{[i,j]}}{\sum_{k=1}^{hw} ReLU(M)_{[i,k]}}$$
(5)

The last step is to matrix product M' and the input z to produce the final output z':

$$z' = M' \cdot z \in \mathbb{R}^{hw \times C} \tag{6}$$

Now z' is the new self-attended output attended by M'. Each logit (row) of z' is a weighted sum of the logits of z whose weights are the corresponding row of M':

$$z'_{i} = \sum_{j=1}^{hw} M'_{[i,j]} z_{j}.$$
(7)

Here, z'_i and z_j refer to a logit (row) of z' and z respectively. Therefore, z' is a set of self-attended logits that captures

(a) The Self-Attention Module Training Phase



Figure 2. Training of the self-attention module and the segmentation network. (a) The module is trained only on the source domain to learn the correct pixel relations via the ground truth labels of the source domain. (b) The source-trained self-attention module is used for the domain adaptive training of the segmentation network. The segmentation network is trained by making z to follow the output of the module, z' or z'' along with self-training scheme. \mathcal{L}_{att} can be computed for both domains.

the information about the inter-pixel correlations based on the similarity between pixels informed by the attention map M'. Note that we use z instead of U(z) as our input to the module because using U(z) requires expensive computation and huge memory when computing M.

Fig. 2 (a) shows the training of the self-attention module. As mentioned earlier, we train the self-attention module only on the source domain because we want the module to learn the domain-invariant inter-pixel correlations and the source domain has the ground truth labels that can guide this correctly. We add the output of the segmentation network z_s to the output of the module z'_s to create $z''_s = z_s + z'_s$. We upsample and softmax z''_s to produce $p''_s = S(U(z''_s))$ and compute the segmentation loss (1) using it. The segmentation network and the self-attention module are jointly trained by minimizing the source segmentation loss. However, when the training is finished, the trained segmentation network is discarded and only the self-attention module is saved to be used later for the next step which is the actual domain adaptive training of the segmentation network.

When computing the source segmentation loss (1), z''_s is used instead of z'_s to make a skip connection which enables the loss signal to directly flow to the segmentation network during back-propagation. If there were no skip connection and the loss were computed on z'_s , the segmentation network would have been trained to produce output that only performs well after the attention module is applied, eventually producing inaccurate z_s . This is not desirable because the attention module has to learn well-represented z_s that can perform well on the source domain in order to capture the correct inter-pixel correlation of the source domain. The skip connection prevents this problem and enables the segmentation network to be trained good enough to perform segmentation correctly on the source domain.

4.2. Domain adaptive training via SAM

In this section, we explain the domain adaptive training of the segmentation network using the SAM that is pretrained on the source domain. The overall training process is shown in Fig. 2 (b). Please note that the segmentation network in Fig. 2 (b) is a different network from the one in Fig. 2 (a). We train another separate segmentation network from scratch and this network is the one to be tested for the performance measurement. Since this step is the actual process of domain adaptation, we train the network for both the source and the target doamin datasets. We input images of both domains into the network and get z_s and z_t respectively. Then, z_s and z_t are used for two different loss functions. One is the commonly used segmentation loss previously introduced in Sec. 3. More specifically, we upsample and apply softmax on z_s and z_t to get p_s and p_t to compute (1) and (3). The other loss function is our proposed selfattention loss. z_s and z_t are given as inputs to the SAM that has already been pre-trained on the source domain. Inform that the pre-trained SAM is frozen and its weights are not updated during the domain adaptive training of the network. The SAM produces z'_s and z'_t which are the self-attended logits. We add z_s and z_t to them respectively and get z_s'' and z_t'' . The proposed self-attention loss is defined as the L1 loss between z and z' or between z and z''.

$$\mathcal{L}_{att}^{S}(x_{s}) = ||z_{s} - z_{s}''||_{1}^{1} \quad \mathcal{L}_{att}^{T}(x_{t}) = ||z_{t} - z_{t}''||_{1}^{1} \quad (8)$$

$$\mathcal{L}_{att}^{S}(x_{s}) = ||z_{s} - z_{s}'||_{1}^{1} \quad \mathcal{L}_{att}^{T}(x_{t}) = ||z_{t} - z_{t}'||_{1}^{1} \quad (9)$$

The self-attention loss is defined in two forms because we empirically observe that the self-attention loss shows different performances for different source datasets. The intention of the self-attention loss is to train the segmentation network to learn the self-attended output which contains information about inter-pixel correlations by forcing z to mimic the output of the module, z' or z''. When computing the self-attention loss, we detach the output of the module, z' and z'', so the loss signal only flows through z during the back-propagation. The overall objective functions are defined as follows:

$$\mathcal{L} = \mathcal{L}_{seg}^{S}(x_{s}) + \lambda \mathcal{L}_{att}^{S}(x_{s}) + \mathcal{L}_{seg}^{T}(x_{t}) + \lambda \mathcal{L}_{att}^{T}(x_{t}).$$
(10)
$$\mathcal{L} = \mathcal{L}_{seg}^{S}(x_{s}) + \mathcal{L}_{seg}^{T}(x_{t}) + \lambda \mathcal{L}_{att}^{T}(x_{t}).$$
(11)

 λ is a hyper-parameter that balances between the segmentation loss and the self-attention loss. We empirically find that for GTA5 \rightarrow Cityscapes task, using (8) and applying selfattention loss for both domains, (10), performs better but for SYNTHIA \rightarrow Cityscapes, using (9) and applying it only on the target domain, (11), shows better results. This will be further studied in the ablation study of Sec. 5.2. In the inference time, prediction is computed solely based on *z* without using the SAM, hence the SAM is not needed.

5. Experiments

5.1. Datasets and Training Details

Datasets. We conduct UDA experiments using two different source domain datasets, GTA5 [19] and SYN-THIA [20] which are photo-realistic synthetic datasets collected from virtual games scenes. The target domain is the Cityscapes [6] which consists of real scenes under driving scenarios. The GTA5 contains 24,966 images and shares 19 classes with the Cityscapes while the SYNTHIA dataset contains 9,400 images and shares 16 classes in common with the Cityscapes. We additionally report the results of 13 common classes for SYNTHIA. The Cityscapes dataset contains 2,975 training images and 500 validation images. We conduct experiments on the validation set following the standard protocols of previous works.

Training Details. The experiments are mainly conducted using DeepLabV2 [3] with ResNet101 [10] backbone. We also conduct experiment using FCN-8s [16] with VGG16 backbone [22] but its results and training details are in the supplementary. We initialize the backbone of the network with ImageNet pre-trained weights. We use the SGD optimizer with initial learning rate of 2.5×10^{-4} and weightdecay of 0.0005. Learning rate is scheduled using 'poly' learning rate policy with a power of 0.9. During the selfattention module training phase, both segmentation network and the SAM are jointly trained by the same optimizer. When domain adaptive training the segmentation network via the self-attention module, we use the pre-trained selfattention module that is saved at the last save point. The network is trained for 120,000 iterations with a batch size of 1 and is tested and saved at every 2,000 iterations. The

Method	GTA5→CS	SYNTHIA→CS
No Pseudo	46.8	39.3
Pseudo-Only	48.9	43.1
Ours (8) & (10)	50.7	43.5
Ours (9) & (10)	50.3	44.4
Ours Only on Target (11)	50.3 (8)	45.1 (9)
No Conv	42.7	44.0
No skip-connection	49.7	44.4

Table 1. Results of Ablation Studies. The numbers are mIoU of 19 classes and 16 classes for GTA5 \rightarrow CS and SYNTHIA \rightarrow CS respectively.

	GTA5→C	CS	SYNTHIA→CS							
Gen	Pseudo-Only	Ours	Pseudo-Only	Ours						
Gen1	48.9	50.7	43.1	45.1						
Gen2	49.7	51.2	45.2	47.9						
Gen3	49.6	51.4	45.7	49.0						
Gen4	49.1	51.7	45.9	49.2						
Gen5	48.8	52.2	45.7	48.5						

Table 2. Results of Iterative Training. The best results are in bold. mIoU 19 and mIoU 16 are used for $GTA5 \rightarrow CS$ and $SYN-THIA \rightarrow CS$ respectively.

hyper-parameter λ is set to 0.1 empirically. We utilize the source images that are style-transferred into Cityscapes by [42] to further close the domain gap at image level.

5.2. Ablation Study

Tab.1 shows our ablation studies. We test our proposed method in different settings and validate which setting works the best for each UDA task. 'No Pseudo' does not employ any self-training via pseudo labels and only trains the network with the style-transferred source domain images. It shows the lowest performance for both UDA tasks. 'Pseudo-Only' refers to applying self-training using the pseudo labels generated by the 'No Pseudo' models. It is trained only by the two loss terms (1) and (3). It clearly shows improved performance on both tasks. 'Ours (8)' and 'Ours (9)' are our proposed method introduced in Sec.4.2. We utilize the same pseudo labels used in 'Pseudo-Only' and additionally employ our proposed self-attention loss to show the efficacy of our method. The performances are significantly improved when our method is applied. However, there are some different tendency between the two UDA tasks. For GTA5→Cityscapes, (8) shows better performance than (9) while (9) results in higher mIoU than (8) on SYNTHIA \rightarrow Cityscapes. Also, when we apply the self-attention loss only on the target domain ('Ours Only on Target'), employing (11), the performance on SYNTHIA improves but it rather decreases on GTA5. We conjecture this is due to the larger domain gap between SYNTHIA and Cityscapes than GTA5 and Cityscapes. (See supplementary for further discussion about it.) Therefore, we utilize (8) and (10) on GTA5 \rightarrow Cityscapes while (9) and (11) on SYN-THIA->Cityscapes for the studies of 'No Conv' and 'No

Task	Method	road	side.	build.	wall	fence	pole	light	sign	vege.	terrain	sky	person	rider	car	truck	bus	train	motor	bike	mIoU
GTA5→CS	Pseudo-Only Ours	90.6 91.0	43.7 46.9	84.6 85.7	37.6 37.6	24.8 25.5	33.8 34.2	39.4 39.4	36.9 37.7	84.6 85.3	38.8 36.3	84.3 85.4	59.7 62.9	26.0 33.9	85.3 85.4	39.5 42.9	50.0 50.9	0.0 0.0	24.6 33.4	45.1 48.1	48.9 50.7
SYNTHIA →CS	Pseudo-Only	72.7	29.2 28.6	80.0 82.2	13.8	0.8	31.3	24.5	28.4 28.9	80.3 81.7	_	82.4 85.0	50.9 56.1	24.7	74.0	_	24.4	_	21.7	49.9 50 1	43.1

Table 3. Effectiveness of 'Ours' by classes. Red and Blue refer to increase and decrease in performance. Bold faced are the rare classes.

Skip-connection'. Each refers to different architecture setting of SAM, 'No Conv' trains the network using a SAM without the 1×1 convolutional layer, thus it does not need any training on the source domain and the attention map is directly computed on z without transformation. 'No Skipconnection' trains the SAM without the skip-connection, hence the source segmentation loss is computed on z' not on z'' during the SAM training phase. Without the skipconnection, the module fails to learn well-represented z_s and it eventually under-performs during the domain adaptive training of the segmentation network. In both SAM settings, we could observe the performance degradation for both UDA tasks. This validates the necessity of the 1×1 convolutional layer and the skip-connection of the SAM.

5.3. Iterative Training

Self-training methods [15, 36, 35, 37] employ iterative training scheme which generates pseudo labels for the target domain from a trained segmentation network and retrains a new network with the generated pseudo labels. Tab. 2 shows the iterative training of 'Pseudo-Only' and 'Ours' on both UDA tasks. 'Gen' stands for a generation which indicates the iteration of the iterative training. For example, Gen1 networks are trained with pseudo labels generated by the 'No Pseudo' models from Tab. 1 and Gen2 networks are trained with pseudo labels generated by Gen1 networks. On GTA5 task, the performance of 'Pseudo-Only' shows the highest at Gen2 and rather decreases after as the generation goes on, while on SYNTHIA task, there is only a marginal improvement between the generations. On the other hand, when training with our proposed self-attention loss, the performance tends to improve better at every generation for both source domains. It implies that our proposed method helps the network not to produce noisy error-prone pseudolabels that can lead to incorrect supervision signals and to generalize better on the target domain by transferring domain invariant inter-pixel correlations from the source domain. Note that the SAM is trained only *once* on the source domain and the same SAM is reused at every generation for the proposed self-attention loss.

5.4. Class-wise comparison between Pseudo-Only and Ours

In Tab. 3 we compare our method with 'Pseudo-Only' by classes to further investigate the effectiveness of our proposed method. As mentioned earlier, the same pseudo labels are utilized for both methods. 'Ours' only additionally employs the proposed self-attention loss along with selftraining. For both tasks, the mIoU is increased after applying our method. Our method particularly contributes to the performance gain on rare classes which scarcely exist in the dataset. Wall, fence, pole, light, sign, terrain, rider, truck, bus, train, motor and bike are known rare classes. On GTA5→Cityscapes, IoU of rider, truck, motor, and bike are significantly improved, besides, other rare classes such as fence, pole, sign and bus also show increased IoU. Similarly on SYNTHIA→Cityscapes, 'Ours' shows better performance than 'Pseudo-Only' on pole, light, sign, rider, bus, motor and bike. Other than rare-classes, our method also leads to significant improvement of majors classes such as sky, person and car on both tasks.

5.5. Comparisons with state-of-the-art methods

Tab. 4 and Tab. 5 show the comparison experiments with other state-of-the-art methods on both datasets. Our method outperforms other methods except for ProDA [36] which achieved a major breakthrough in the UDA for semantic segmentation. Since our proposed method can be easily applied to other methods by simply adding our self-attention loss term to the overall objective function, we combine our method with ProDA by employing our self-attention loss during the stage3 training phase of ProDA. After applying our method on ProDA, we could observe an additional performance gain for both UDA tasks. On GTA5→Cityscapes, 12 out of 19 classes show increased IoU, especially, the performance on rare classes such as fence, pole, terrain, truck and bus are highly improved. Similar results can be found on SYNTHIA-Cityscapes as well, mIoU is improved along with IoU of some rare-classes. Also, it is noteworthy that when excluding ProDA, there is a more than 1% gap between 'Ours' (49.2%) and the second-best method (47.9%) on SYNTHIA task. We empirically find that $\lambda =$ 0.005 works the best when combining our method with ProDA. On both UDA benchmarks, our method achieves comparable performance to existing approaches. Moreover, we show that our method can be easily applied to other method and validate its effectiveness by additionally employing it on leading SOTA method and setting a new SOTA performance for both UDA tasks.

Method	road side	. build.	wall	fence	pole	light	sign	vege.	terrain	sky	person	rider	car	truck	bus	train	motor	bike	mIoU
PLCA [13]	84.0 30.4	82.4	35.3	24.8	32.2	36.8	24.5	85.5	37.2	78.6	66.9	32.8	85.5	40.4	48.0	8.8	29.8	41.8	47.7
BDL [15]	91.0 44.7	84.2	34.6	27.6	30.2	36.0	36.0	85.0	43.6	83.0	58.6	31.6	83.3	35.3	49.7	3.3	28.8	35.6	48.5
CrCDA [12]	92.4 55.3	8 82.3	31.2	29.1	32.5	33.2	35.6	83.5	34.8	84.2	58.9	32.2	84.7	40.6	46.1	2.1	31.1	32.7	48.6
SIM [31]	90.6 44.7	84.8	34.3	28.7	31.6	35.0	37.6	84.7	43.3	85.3	57.0	31.5	83.8	42.6	48.5	1.9	30.4	39.0	49.2
CADA [34]	91.3 46.0	84.5	34.4	29.7	32.6	35.8	36.4	84.5	43.2	83.0	60.0	32.2	83.2	35.0	46.7	0.0	33.7	42.2	49.2
Label-driven[33]	90.8 41.4	84.7	35.1	27.5	31.2	38.0	32.8	85.6	42.1	84.9	59.6	34.4	85.0	42.8	52.7	3.4	30.9	38.1	49.5
MaxCos [5]	92.6 54.0	85.4	35.0	26.0	32.4	41.2	29.7	85.1	40.9	85.4	62.6	34.7	85.7	35.6	50.8	2.4	31.0	34.0	49.7
Kim et al. [14]	92.9 55.0	85.3	34.2	31.1	34.9	40.7	34.0	85.2	40.1	87.1	61.0	31.1	82.5	32.3	42.9	0.3	36.4	46.1	50.2
CAG_UDA [37]	90.4 51.6	5 83.8	34.2	27.8	38.4	25.3	48.4	85.4	38.2	78.1	58.6	34.6	84.7	21.9	42.7	41.1	29.3	37.2	50.2
Seg-Uncertainty [39]	90.4 31.2	2 85.1	36.9	25.6	37.5	48.8	48.5	85.3	34.8	81.1	64.4	36.8	86.3	34.9	52.2	1.7	29.0	44.6	50.3
FDA-MBT [35]	92.5 53.3	8 82.4	26.5	27.6	36.4	40.6	38.9	82.3	39.8	78.0	62.6	34.4	84.9	34.1	53.1	16.9	27.7	46.4	50.5
TPLD [21]	94.2 60.5	5 82.8	36.6	16.6	39.3	29.0	25.5	85.6	44.9	84.4	60.6	27.4	84.1	37.0	47.0	31.2	36.1	50.3	51.2
MetaCorrection [9]	92.8 58.1	86.2	39.7	33.1	36.3	42.0	38.6	85.5	37.8	87.6	62.8	31.7	84.8	35.7	50.3	2.0	36.8	48.0	52.1
Ours	90.8 47.2	86.8	41.5	29.4	35.7	42.4	37.4	86.0	42.1	88.3	63.7	35.6	85.1	43.8	54.6	0.0	33.6	47.8	52.2
ProDA [36]	87.8 56.0) 79.7	46.3	44.8	45.6	53.5	53.5	88.6	45.2	82.1	70.7	39.2	88.8	45.5	59.4	1.0	48.9	56.4	57.5
ProDA+Ours	89.3 56.4	81.1	46.1	46.2	47.0	52.9	54.1	88.7	48.3	82.5	70.4	39.1	89.3	51.8	61.1	0.1	44.4	55.0	58.1

Table 4. Comparison results with other methods on GTA5→Cityscapes. The numbers in bold are the best score for each column.

Method	road	side.	build.	wall	fence	pole	light	sign	vege.	sky	person	rider	car	bus	motor	bike	mIoU*	mIoU
Kim et al. [14]	92.6	53.2	79.2	—	_	—	1.6	7.5	78.6	84.4	52.6	20.0	82.1	34.8	14.6	39.4	49.3	
CrCDA [12]	86.2	44.9	79.5	8.3	0.7	27.8	9.4	11.8	78.6	86.5	57.2	26.1	76.8	39.9	21.5	32.1	50.0	42.9
BDL [15]	86.0	46.7	80.3	—	—	—	14.1	11.6	79.2	81.3	54.1	27.9	73.7	42.2	25.7	45.3	51.4	
CAG_UDA [37]	84.7	40.8	81.7	7.8	0.0	35.1	13.3	22.7	84.5	77.6	64.2	27.8	80.9	19.7	22.7	48.3	51.5	44.5
SIM [31]	83.0	44.0	80.3	—	—	—	17.1	15.8	80.5	81.8	59.9	33.1	70.2	37.3	28.5	45.8	52.1	_
CADA [34]	82.5	42.2	81.3	_	—	_	18.3	15.9	80.6	83.5	61.4	33.2	72.9	39.3	26.6	43.9	52.4	
FDA-MBT [35]	79.3	35.0	73.2	—	—	—	19.9	24.0	61.7	82.6	61.4	31.1	83.9	40.8	38.4	51.1	52.5	_
MetaCorrection [9]	92.6	52.7	81.3	8.9	2.4	28.1	13.0	7.3	83.5	85.0	60.1	19.7	84.8	37.2	21.5	43.9	52.5	45.1
Label-driven[33]	85.1	44.5	81.0	—	—	—	16.4	15.2	80.1	84.8	59.4	31.9	73.2	41.0	32.6	44.7	53.1	_
MaxCos [5]	88.3	47.3	80.1	_	—	_	21.6	20.2	79.6	82.1	59.0	28.2	82.0	39.2	17.3	46.7	53.2	
PLCA[13]	82.6	29.0	81.0	11.2	0.2	33.6	24.9	18.3	82.8	82.3	62.1	26.5	85.6	48.9	26.8	52.2	54.0	46.8
TPLD [21]	80.9	44.3	82.2	19.9	0.3	40.6	20.5	30.1	77.2	80.9	60.6	25.5	84.8	41.1	24.7	43.7	53.5	47.3
Seg-Uncertainty [39]	87.6	41.9	83.1	14.7	1.7	36.2	31.3	19.9	81.6	80.6	63.0	21.8	86.2	40.7	23.6	53.1	54.9	47.9
Ours	77.5	32.3	82.6	25.5	1.9	34.6	33.6	32.4	81.7	85.1	63.8	31.8	82.3	35.2	31.9	54.6	55.7	49.2
ProDA [36]	87.8	45.7	84.6	37.1	0.6	44.0	54.6	37.0	88.1	84.4	74.2	24.3	88.2	51.1	40.5	45.6	62.0	55.5
ProDA [36] + Ours	87.3	45.4	85.3	37.3	0.3	46.2	54.2	39.0	88.3	84.6	73.7	22.9	88.5	59.8	41.8	44.1	62.7	56.2

Table 5. Comparison results with other methods on SYNTHIA \rightarrow Cityscapes. The numbers in bold are the best score for each column. mIoU* and mIoU denote mIoU of 13 classes and 16 classes respectively.

5.6. Qualitative Analysis

Fig. 3 shows attention and prediction visualization of 'Pseudo-Only' and 'Ours'. 'Pseudo-Only' and 'Ours' are networks tested in Tab. 3. The attention is visualized as the similarity between the pixel with the highest confidence of the class we want to measure and all the other pixels in the image. The most confident pixel is marked as a blue dot in the visualization. The redder color means higher similarity while the bluer color refers to lower similarity. If the network is trained well, the most confident pixel should have high similarity only with the pixels belonging to the same class. We additionally visualize the results after applying the SAM using z'' and z' ('Ours-z''' & 'Ours-z'') for GTA5 and SYNTHIA respectively.

When using only the self-training ('Pseudo-Only'), the most confident pixel has high similarity even with irrelevant pixels of different classes, while 'Ours' is more attended well with the relevant pixels corresponding to the same class. The results of SYNTHIA are more noisy due to the larger domain gap. 'Ours-z''' and 'Ours-z'' look a little more noisy than 'Ours' but clearer and more distinct than 'Pseudo-Only'. They give us hint about what knowledge is being transferred via our proposed self-attention loss. We conjecture that their little noisiness provides extra information about inter-pixel correlations to the network. In the prediction visualization, the 3rd and the 4th columns correspond to GTA5 task while the 5th and the 6th columns refer to SYNTHIA task. When comparing only the prediction results, it is hard to see the difference between the two methods, even though 'Ours' look a little more accurate. However, when we visualize the attention by classes and compare them as mentioned above, we can clearly observe a distinct difference between the two methods. It implies that even though their prediction visualization might seem similar, their predicted logits have very different distribution which leads to better performance for 'Ours'. These re-



Figure 3. The attention and prediction visualization. Images above the blue dashed line are attention visualization and ones under the line are prediction visualization. In the attention visualization, the images on the left side of the red dashed line are the results of $GTA5 \rightarrow Cityscapes$ while ones on the right are SYNTHIA \rightarrow Cityscapes. Each row of attention visualization refers to 'pole', 'light', 'sign', 'vege.' and 'person' from top to bottom. One image is used to fairly show the difference between different classes and methods.



Figure 4. Entropy analysis of $GTA5 \rightarrow Cityscapses$ (Left) and $SYNTHIA \rightarrow Cityscapses$ (right).

sults are not cherry-picked and more results using different images can be found in the supplementary.

5.7. Quantitative Analysis

Fig. 4 shows entropy analysis of three different networks on the two UDA tasks. Every 2,000 iteration, we measure the entropy of all the 500 validation images of Cityscapes. For each image, we divide the pixels into two groups, correctly and incorrectly classified, and compute the entropy for each group as follows:

$$Ent(x_t) = \frac{1}{\#\text{pixel}} \left(-\sum_{i=1}^{\#\text{pixel}} \sum_{c=1}^{C} p_t^{(i,c)} \log(p_t^{(i,c)})\right) \quad (12)$$

We plot the average entropy of the 500 images for each pixel group separately in the figure. The solid line is the entropy of correct pixels and the dashed line refers to the incorrect pixels. Each red, blue and green line indicates 'Ours-(8)', 'Ours-(9)' and 'Pseudo-Only' networks respectively. For both benchmarks and all three networks, incorrect pixels have higher entropy than correct pixels. It is an interesting observation that 'Pseudo-Only' has the lowest entropy while 'Ours-(8)' and 'Ours-(9)' show higher entropy than 'Pseudo-Only' throughout the training. Having higher entropy means that the class distribution of a pixel is not biased towards a certain class (a highly peaky distribution) but rather has probability for some of other classes as well. We conjecture this is due to our self-attention loss providing extra knowledge about inter-pixel correlations which prevents the network from being highly biased and over-confident.

6. Conclusion

In this paper, we propose a method of transferring domain invariant inter-pixel correlations from the source domain to the target domain by utilizing a self-attention module that is pre-trained on the source domain. Our method is very simple and does not require complex and heuristic algorithm (only one hyper-parameter to tune, λ) yet very effective. Our method supports self-training by overcoming the defects of noisy pseudo labels with extra knowledge provided by the SAM. Also, we show that our proposed method can be easily used in addition to other existing methods to further enhance the performance. The proposed method sets a new SOTA scores on both UDA benchmarks when combined with the recent SOTA method [36].

7. Acknowledgment

This work was supported by NRF grant (2021R1A2C3006659) and IITP grants (No.2021-0-01343, No.2022-0-00320), all of which are funded by Korean Government.

References

- Abien Fred Agarap. Deep learning using rectified linear units (relu). arXiv preprint arXiv:1803.08375, 2018.
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-toend object detection with transformers. In *European Conference on Computer Vision (ECCV)*, 2020.
- [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. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [4] Jianpeng Cheng, Li Dong, and Mirella Lapata. Long shortterm memory-networks for machine reading. arXiv preprint arXiv:1601.06733, 2016.
- [5] Inseop Chung, Daesik Kim, and Nojun Kwak. Maximizing cosine similarity between spatial features for unsupervised domain adaptation in semantic segmentation. *arXiv preprint arXiv:2102.13002*, 2021.
- [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 Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [8] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. In *Journal of Machine Learning Research, vol. 17*, pages 1–35, 2016.
- [9] Xiaoqing Guo, Chen Yang, Baopu Li, and Yixuan Yuan. Metacorrection: Domain-aware meta loss correction for unsupervised domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [11] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning*, 2018.
- [12] Jiaxing Huang, Shijian Lu, Dayan Guan, and Xiaobing Zhang. Contextual-relation consistent domain adaptation for semantic segmentation. In *European Conference on Computer Vision*, pages 705–722. Springer, 2020.
- [13] Guoliang Kang, Yunchao Wei, Yi Yang, Yueting Zhuang, and Alexander G Hauptmann. Pixel-level cycle association: A

new perspective for domain adaptive semantic segmentation. arXiv preprint arXiv:2011.00147, 2020.

- [14] Myeongjin Kim and Hyeran Byun. Learning texture invariant representation for domain adaptation of semantic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12975– 12984, 2020.
- [15] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vi*sion and Pattern Recognition, pages 6936–6945, 2019.
- [16] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [17] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3764–3773, 2020.
- [18] Ankur P Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. A decomposable attention model for natural language inference. arXiv preprint arXiv:1606.01933, 2016.
- [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, *European Conference on Computer Vision* (ECCV), volume 9906 of LNCS, pages 102–118. Springer International Publishing, 2016.
- [20] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio Lopez. The SYNTHIA Dataset: A large collection of synthetic images for semantic segmentation of urban scenes. 2016.
- [21] Inkyu Shin, Sanghyun Woo, Fei Pan, and In So Kweon. Twophase pseudo label densification for self-training based domain adaptation. In *European conference on computer vision*, pages 532–548. Springer, 2020.
- [22] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
- [23] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018.
- [24] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [25] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Advances in Neural Information Processing Systems, 2017.
- [26] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In

Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2517–2526, 2019.

- [27] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Mathieu Cord, and Patrick Pérez. Dada: Depth-aware domain adaptation in semantic segmentation. In *ICCV*, 2019.
- [28] Haoran Wang, Tong Shen, Wei Zhang, Ling-Yu Duan, and Tao Mei. Classes matter: A fine-grained adversarial approach to cross-domain semantic segmentation. In *European Conference on Computer Vision*, pages 642–659. Springer, 2020.
- [29] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2018.
- [30] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12275–12284, 2020.
- [31] Zhonghao Wang, Mo Yu, Yunchao Wei, Rogerio Feris, Jinjun Xiong, Wen-mei Hwu, Thomas S Huang, and Honghui Shi. Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12635–12644, 2020.
- [32] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In European Conference on Computer Vision (ECCV), 2018.
- [33] 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.
- [34] Jinyu Yang, Weizhi An, Chaochao Yan, Peilin Zhao, and Junzhou Huang. Context-aware domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 514–524, 2021.
- [35] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4085–4095, 2020.
- [36] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12414–12424, 2021.
- [37] Qiming Zhang, Jing Zhang, Wei Liu, and Dacheng Tao. Category anchor-guided unsupervised domain adaptation for semantic segmentation. In Advances in Neural Information Processing Systems, pages 433–443, 2019.
- [38] Sicheng Zhao, Xiangyu Yue, Shanghang Zhang, Bo Li, Han Zhao, Bichen Wu, Ravi Krishna, Joseph E Gonzalez, Alberto L Sangiovanni-Vincentelli, Sanjit A Seshia, et al. A review of single-source deep unsupervised visual domain adaptation. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.

- [39] Zhedong Zheng and Yi Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. arXiv preprint arXiv:2003.03773, 2020.
- [40] Zhedong Zheng and Yi Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision*, 129(4):1106–1120, 2021.
- [41] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.
- [42] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycleconsistent adversarial networks. In *Computer Vision (ICCV)*, 2017 IEEE International Conference on, 2017.
- [43] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289– 305, 2018.
- [44] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 5982–5991, 2019.