

Cross-Domain Grouping and Alignment for Domain Adaptive Semantic Segmentation

Minsu Kim¹, Sunghun Jung¹, Seungryong Kim², Jungin Park¹, Ig-Jae Kim³, Kwanghoon Sohn¹

¹ Yonsei University ² Korea University

³ Korea Institute of Science and Technology (KIST)

{minkim320, sunghunjung, newrun, khsohn}@yonsei.ac.kr, seungryong_kim@korea.ac.kr, drjay@kist.re.kr

Abstract

Existing techniques to adapt semantic segmentation networks across source and target domains within deep convolutional neural networks (CNNs) deal with all the samples from the two domains in a global or category-aware manner. They do not consider an inter-class variation within the target domain itself or estimated category, providing the limitation to encode the domains having a multi-modal data distribution. To overcome this limitation, we introduce a learnable clustering module, and a novel domain adaptation framework, called cross-domain grouping and alignment. To cluster the samples across domains with an aim to maximize the domain alignment without forgetting precise segmentation ability on the source domain, we present two loss functions, in particular, for encouraging semantic consistency and orthogonality among the clusters. We also present a loss so as to solve a class imbalance problem, which is the other limitation of the previous methods. Our experiments show that our method consistently boosts the adaptation performance in semantic segmentation, outperforming the state-of-the-arts on various domain adaptation settings.

Introduction

Semantic segmentation aims at densely assigning semantic category label to each pixel given an image. Though the remarkable progresses have been dominated by deep neural networks trained on large-scale labeled dataset (Chen et al. 2017a). The segmentation model trained on the labeled data in source domain usually cannot generalize well to the unseen data in target domain. For example, the model trained on the data from one city or computer-generated scene (Richter et al. 2016; Ros et al. 2016) may fail to yield accurate pixel-level predictions for the scenes of another city or real scene. The main reason lies in the different data distribution between such source and target domains, typically known as domain discrepancy (Shimodaira 2000).

To address this issue, domain adaptive semantic segmentation methods have been proposed in which they align data distribution between the source and target domains by adopting a domain discriminator (Hoffman et al. 2016; Tsai et al. 2018). Formally, these methods aim to minimize an adversarial loss (Goodfellow et al. 2014) to reduce the domain discrepancy at image-level (Wu et al. 2018; Hoffman

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

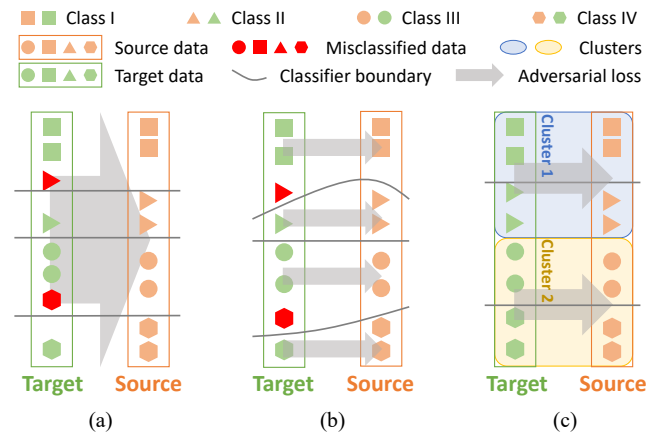


Figure 1: Illustration of cross-domain grouping and alignment : Conventional methods aim to reduce the domain discrepancy between source and target domains through (a) global and (b) category-level domain alignment, without taking into account the *inter-class* variation or rely solely on the category classifier. (c) We propose to replace this category classifier with an intermediate cross-domain grouping module to align each group separately (best view in color).

et al. 2018; Chang et al. 2019), feature-level (Hoffman et al. 2016), and category probability-level (Zou et al. 2018; Li, Yuan, and Vasconcelos 2019; Tsai et al. 2018) distributions without forgetting semantic segmentation ability on the source domain. However, their accuracy is still limited when aligning multi-modal data distribution (Arora et al. 2017), which cannot guarantee that the target samples from different categories are properly separated as in Fig. 1 (a).

To tackle this limitation, category-level domain adaptation methods (Chen et al. 2017b; Du et al. 2019) have been proposed for semantic segmentation in which they minimize the class-specific domain discrepancy across the source and target domains. Together with supervision from the source domain, this enforces the segmentation network to learn discriminative representation for different classes on both domains. They utilize a category classifier trained on the source domain to generate *pseudo* class labels on the target domain. It results in inaccurate labels for domain adaptation that misleads the domain alignment and accumulates errors as in Fig.

1 (b). It also contains a class imbalance problem (Zou et al. 2018), where the network works well for majority categories with a large number of pixels (e.g. road and building), while not suitable for minority categories with a small number of pixels (e.g. traffic sign).

To overcome this limitation, we present cross-domain grouping and alignment for domain adaptive semantic segmentation. As illustrated in Fig. 1 (c), the key idea of our method is to apply an intermediate grouping module to replace the category classifier, allowing to align samples of source and target domains at each group to be similar without using error-prone category classifier. To make the grouping module help with domain adaptation, we propose several losses in a manner that the category distribution of each group between different domains should be consistent, while the category distribution of different groups in the same domain should be orthogonal. Furthermore, we present a group-level class equivalence scheme in order to align all the categories regardless of the number of pixels. The proposed method is extensively evaluated through an ablation study and comparison with state-of-the-art methods on various domain adaptive semantic segmentation benchmarks, including GTA5 \rightarrow Cityscapes and SYNTHIA \rightarrow Cityscapes.

Related Work

Semantic Segmentation

Numerous methods have been proposed to assign class labels in pixel level for input images. Long et al. (2015) first transformed a classification convolutional neural network (CNN) (Krizhevsky, Sutskever, and Hinton 2012; Karen Simonyan and Zisserman 2015; He et al. 2016) to a fully-convolutional network (FCN) for semantic segmentation. Following the line of FCN-based methods, several methods utilized dilated convolutions to enlarge the receptive field (Yu and Koltun 2015) and reason about spatial relationship (Chen et al. 2017a). Recently, Zhao et al. (2017) presented pyramid pooling module to encode the global and local context. Although these methods yielded impressive results in semantic segmentation, they still relied on large datasets with dense pixel-level class labels, which is expensive and laborious. An alternative is to utilize synthetic data (Richter et al. 2016; Ros et al. 2016) which can make unlimited amounts of labels available. Nevertheless, synthetic data still suffer from a substantially different data distribution from real data, which results in a dramatic performance drop when applying the trained model to real scenes.

Domain Adaptive Semantic Segmentation

Due to the obvious mismatch between synthetic and real data, unsupervised domain adaptation (UDA) is studied to minimize the domain discrepancy by aligning the feature distribution between source and target data. As a pioneering work, Ganin et al. (2015) introduced the domain adversarial network to transfer the feature distribution, and Tzeng et al. (2017) proposed adversarial discriminative alignment.

For pixel-level classification, numerous approaches (Wu et al. 2018; Hoffman et al. 2018; Chang et al. 2019) utilized

image-level adaptation methods which translate source image to have the texture appearance of target image, while preserving the structure information of the source image for adapting cross-domain knowledge. In contrast, several methods (Zou et al. 2018; Li, Yuan, and Vasconcelos 2019; Li et al. 2020) adopted the iterative self-training approach to alternatively select unlabelled target samples with higher class probability and utilized them as a pseudo ground-truth. The feature-level adaptation methods align the intermediate feature distribution via adversarial framework. Hoffman et al. (2016) introduced a feature-level adaptation method to align the intermediate feature distribution for the global and local alignment. Tsai et al. (2018) adopted output-level adaptation for structured output space, since it contains similar spatial structure with semantic segmentation. However, these methods aim to align overall data distribution without taking into account the inter-class variation.

To solve this problem, several methods (Chen et al. 2017b; Du et al. 2019) introduced category-level adversarial learning to align the data distributions independently for each class. Similarly, other works (Tsai et al. 2019; Huang et al. 2020) discovered patch-level adaptation methods by using multiple modes of patch-wise output distribution to differentiate the feature representation of patches. However, inaccurate domain alignment occurs because these methods rely heavily on category or patch classifiers trained in the source domain. The most similar to our work is Wang et al. (2020), which group the category classes into several groups for domain adaptive semantic segmentation. While they divide stuff and things (i.e. disconnected regions), our cross-domain grouping module divides the categories into multiple groups that the grouping network and segmentation network can be trained in a joint and boosting manner.

Unsupervised Deep Clustering

A variety of approaches have applied deep clustering algorithms that simultaneously discover groups in training data and perform representation learning. Chang et al. (2017) proposed to cast the clustering problem into pairwise classification using CNN. Caron et al. (2018) proposed a learning procedure that alternates between clustered images in the representation space and trains a model that assigns images to their clusters. Other approaches (Joulin, Bach, and Ponce 2012; Tao et al. 2017) localized the salient and common objects by clustering pixels in multiple images. Similarly, Collins et al. (2018) proposed deep feature factorization (DFF) to group the common part segments between images through non-negative matrix factorization (NMF) (Ding, He, and Simon 2005) on CNN features. This paper follows such a strategy to group semantic consistent data representation across the source and target domains.

Proposed Method

Problem Statement and Overview

Let us denote the source and target images as I_S, I_T , where only the source data is annotated with per-pixel semantic categories as Y_S . We seek to train a semantic segmentation network G , which outputs pixel-wise class probabil-

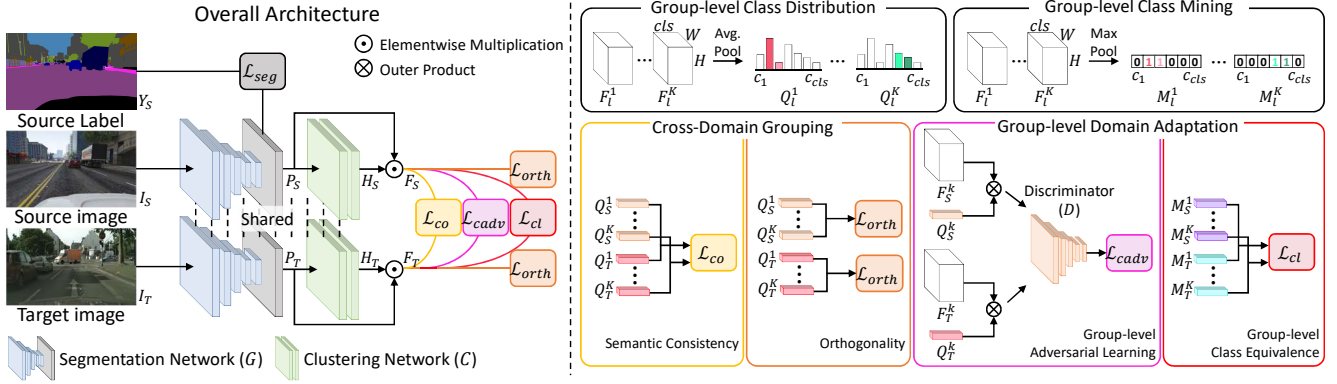


Figure 2: Overview of our method. Images from the source and target domains are passed through segmentation network G . We decompose the data distribution of source and target domains into a set of K sub-spaces with cross-domain grouping network C . Then discriminator D distinguishes whether the data distribution for each sub-space is from the source or target domain.

ity P_S, P_T on both source and target domains reliably, with height h , width w , and the number of classes cls , respectively. Our goal is to train the segmentation network that yields to align probability distribution of the source and target domains P_S and P_T so that the network G can correctly predict the pixel-level labels even for the target data I_T , following recent study (Tsai et al. 2018) of adaptation in the output probability space, which shows better performance than adaptation in the intermediate feature space.

Conventionally two types of domain adaptation approaches have been proposed: *global* domain adaptation and *category-level* domain adaptation. The former aims to align the global domain differences, while the latter aims to minimize class specific domain discrepancy for each category. However, *global* domain adaptation does not take into account the inter-class variations, and *category-level* domain adaptation rely solely on category classifier. To this end, we propose a novel method by clustering the samples as K groups across the source and target domains. Concretely, we cluster the probability distribution into K groups using cross-domain grouping module, followed by *group-level* domain alignment. By setting K greater than 1, domain alignment of complicated data distribution can be solved by an alignment of K simple data distributions which is the challenge in *global* domain adaptation. By setting K less than cls , the domain misalignment in *category-level* can be mitigated without using a category classifier trained in the source domain. In the following, we introduce our overall network architecture (Section), several constraints for cross-domain grouping (Section), and cross-domain alignment (Section).

Network Architecture

Fig. 2 illustrates our overall framework. Our network consists of three major components: 1) the semantic segmentation network G , 2) the cross-domain grouping module C to cluster sub-spaces based on the output probability distribution, and 3) the discriminator D for group-level domain adaptation. In the following sections, we denote source and target domains as $l \in \{S, T\}$ unless otherwise stated.

Segmentation network. Following the works (Tsai et al. 2018; Li, Yuan, and Vasconcelos 2019; Wang et al. 2020), we exploit DeepLab-V2 (2017a) with ResNet-101 (2016) pre-trained on ImageNet (2009) dataset. The source and target images I_l are fed into the segmentation network G , outputting pixel-wise class probability distribution $P_l = G(I_l)$. Note that P_l is extracted from the segmentation network before applying a softmax layer with same resolution as the input using bilinear interpolation, similar to Tsai et al. (2018).

Cross-Domain grouping network. Our cross-domain grouping network C is formulated as two convolutions. We design each convolution with 1×1 kernel and group mapping function. The first convolution produces 64-channel feature, followed by ReLU and batch normalization. The second convolution produces K grouping scores, followed by softmax function to output group probability $H_l^k = C(P_l)$. We then apply element-wise multiplication between H_l^k and each channel dimension in P_l , obtaining group-specific feature F_l^k . The cross-domain grouping network can be easily replaced with other learnable clustering methods.

Discriminator. For group-level domain alignment, we fed F_l^k into the discriminator D . Following Li et al. (2019), we set the discriminator using five 4×4 convolutional layers of stride 2, where the number of channels is $\{64, 128, 256, 512, 1\}$ to form the network. We use a leaky ReLU (2013) parameterized by 0.2 which is utilized for each convolutional layer except the last one.

Losses for Cross-Domain Grouping

Perhaps one of the most straightforward ways of grouping is to utilize existing clustering methods, *e.g.* k-means (Coates and Ng 2012) or non-negative matrix factorization (NMF) (Collins, Achanta, and Susstrunk 2018). These strategies, however, are not learnable, and thus, they cannot weave the advantages of category-level domain information. Unlike these, we present a learnable clustering module with two loss functions to take advantage of the category-level domain adaptation methods. We discuss in more detail about the effectiveness of our grouping

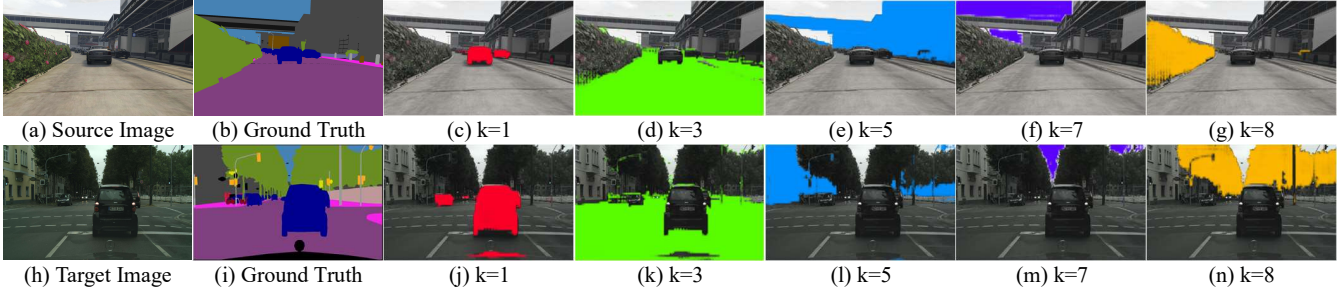


Figure 3: Visualization of cross-domain grouping on the source (first row) and target (second row) image with $K = 8$. (From left to right) Input image and clustering results. Note that color represents the K different sub-space.

compared to non-learnable models (Coates and Ng 2012; Collins, Achanta, and Susstrunk 2018) in Section . In the following, we present each loss function in detail.

Semantic consistency. Our first insight about grouping is that the category distribution of each group between the source and target domains has to be consistent so that the clustered group can benefit from the category-level domain adaptation method. To this end, we first estimate the class distribution Q_l^k by using average pooling layer on each group-level feature F_l^k such that Q_l^k , where each elements in $Q_l^k = [q_1^k, \dots, q_{cls}^k]$ indicates the probability distribution of containing a particular categories at k^{th} group . We then encourage a semantic consistency among class distribution by utilizing l2-norm $\|\cdot\|_2$ as follows:

$$\mathcal{L}_{co}(G, C) = \sum_{k \in \{1, \dots, K\}} \|Q_S^k - Q_T^k\|^2. \quad (1)$$

Minimizing loss (1) has two desirable effects. First, it encourages the difference of each class distribution of group to be similar, and it also provides the supervisory signals for aligning the probability distribution of group-level features.

Orthogonality. The semantic consistency constraint in (1) encourages the class distribution of group across the source and target domains to be consistent. This, however, does not guarantee that class distribution is different for each group. In other words, we cannot divide the multi-modal complex distribution into several simple distributions. To this end, we draw the second insight by introducing orthogonality constraint such that, any two class distribution $Q_l^{j_1}$ and $Q_l^{j_2}$, should be orthogonal each other. It can be realized that their cosine similarity (2) is 0 since Q_l^k are non-negative value. We define the cosine similarity with l2-norm $\|\cdot\|_2$ as follows:

$$\cos(Q_l^{j_1}, Q_l^{j_2}) = \frac{Q_l^{j_1} \cdot Q_l^{j_2}}{\|Q_l^{j_1}\|_2 \|Q_l^{j_2}\|_2}, \quad j_1, j_2 \in \{1, \dots, K\}. \quad (2)$$

We then formulate an orthogonal loss for training such that

$$\mathcal{L}_{orth}(G, C) = \sum_l \sum_{j_1, j_2} \cos(Q_l^{j_1}, Q_l^{j_2}), \quad (3)$$

where we apply a loss function on each domain $l \in \{S, T\}$. By forcing the cross-domain grouping module C to make

each group to be orthogonal, it can divide a multi-modal complex distribution into the K simple class distributions.

Losses for Cross-Domain Alignment

In this section, we present a group-level adversarial learning framework as an alternative to *global* domain adaptation and *category-level* domain adaptation.

Group-level alignment. To achieve group-level domain alignment, a straight forward method is to use K independent discriminators, similar to conventional category-level domain alignment methods (Chen et al. 2017b; Du et al. 2019). However, we simultaneously update grouping module C while training the overall network, thus cluster assignment may not be consistent at each training iteration. To this end, we adopt conditional adversarial learning framework following (Long et al. 2018), by combining group-level feature F_l^k with Q_l^k as a condition as follows:

$$\mathcal{L}_{cadv}(G, C, D) = - \sum_k [\log(D(F_S^k \otimes Q_S^k))] - \sum_k [\log(1 - D(F_T^k \otimes Q_T^k))], \quad (4)$$

where \otimes represent outer product operation. Note that using group-level feature F_l^k only as input to the discriminator is equivalent to global alignment, while we give a condition by using cross-covariance between F_l^k and Q_l^k as input. This leads to discriminative domain alignment according to the different groups.

Group-level class equivalence. For group-level adversarial learning, the existence of particular classes across different domains is desirable. However, since the number of pixels for particular classes are dominant in each image, it can cause class imbalance problem. Thus adaptation model tends to be biased towards majority classes and ignore minority classes (Zou et al. 2018). To alleviate this, we propose group level class equivalence following Zhao et al. (2018). We first apply max pooling layer for each group level feature M_l^k such that each element of $M_l^k = [m_{l,1}^k, \dots, m_{l,cls}^k]$ is a maximum score for each category corresponding to group k . We then utilize maximum classification score in the source domain m_S^k as a pseudo-label, where we aim to train maximum classification score in target domain m_T^k to be similar.

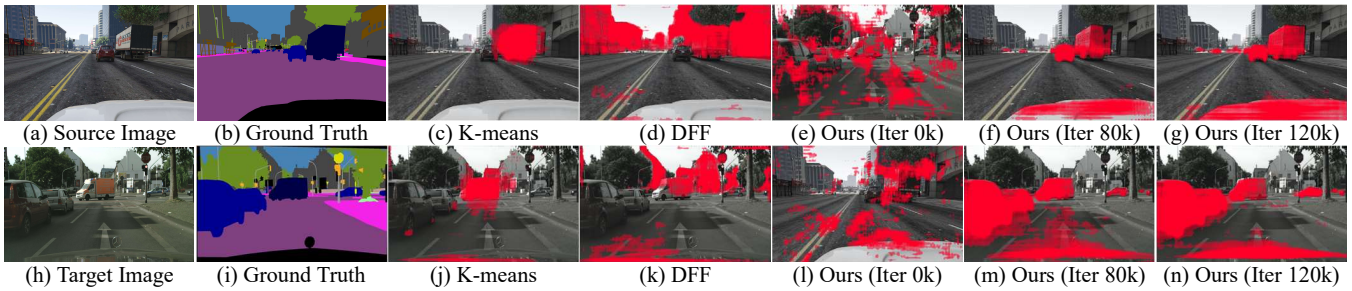


Figure 4: Visualization of cross-domain grouping result of (a) source and (c) target image with GT classes as corresponding colors (b,i) via (c,j) K-means, (d,f) DFF, (e,l) Ours (Iter 0k), (f,m) Ours (Iter 80k) and (g,n) Ours (Iter 120k). Compared to non-learnable model, our method can better capture semantic consistent objects across source and target domains.

To this end, we apply multi-class binary cross-entropy loss for each class as follows :

$$\mathcal{L}_{cl}(G, C) = - \sum_k \sum_u [m_{S,u}^k \geq \tau] \log(m_{T,u}^k), \quad (5)$$

where Iverson bracket indicator function $[\cdot]$ evaluates to 1 when it is true and 0 otherwise, and $u \in \{1, \dots, cls\}$ denotes category. Note that we merely exclude too low probability with the threshold parameter τ .

Training

The overall loss function of our approach can be written as

$$L(G, C, D) = L_{seg}(G) + \lambda_{co}L_{co}(G, C) + \lambda_{cl}L_{cl}(G, C) + \lambda_{orth}L_{orth}(G, C) + \lambda_{cadv}L_{cadv}(G, C, D), \quad (6)$$

where L_{seg} is the supervised cross-entropy loss for semantic segmentation network on the source data, and $\lambda_{co}, \lambda_{orth}, \lambda_{cadv}$ and λ_{cl} are balancing parameter for different losses. We then solve the following minmax problem for optimizing G, C and D .

$$\min_{G,C} \max_D L(G, C, D). \quad (7)$$

Experiments

Experimental Setting

Implementation details. The proposed method was implemented in PyTorch library (Paszke et al. 2017) and simulated on a PC with a single RTX Titan GPU. We utilize BDL (Li, Yuan, and Vasconcelos 2019) as our baseline model following conventional work (Wang et al. 2020), including self-supervised learning and image transferring framework. To train the segmentation network, we utilize stochastic gradient descent (SGD) (1998), where the learning rate is set to 2.5×10^{-4} . For grouping network, we utilize SGD, with learning rate as 1×10^{-3} . Both learning rates decreased with “poly” learning rate policy with power fixed to 0.9 and momentum as 0.9. For discriminator training, we use Adam (2014) optimizer with an initial learning rate 1×10^{-4} . We jointly train our segmentation network, grouping network, and discriminator using (7) for a total of 120k iterations. We randomly paired source and target images in each

iteration. Through the cross-validation using grid-search in log-scale, we set the hyper-parameters $\lambda_{co}, \lambda_{orth}, \lambda_{cadv}, \lambda_{cl}$ and τ as 0.001, 0.001, 0.001, 0.0001 and 0.05, respectively.

Datasets. For experiments, we use the GTA5 (Richter et al. 2016) and SYNTHIA (Ros et al. 2016) as source dataset. GTA5 dataset (Richter et al. 2016) contains 24,966 images with 1914×1052 resolution. We resize images to 1280×760 following other work (Tsai et al. 2018). For SYNTHIA (Ros et al. 2016), we use SYNTHIA-RAND-CITYSCAPES dataset with 9,400 images with 1280×760 resolution. We use Cityscapes (Cordts et al. 2016) as target dataset, which consists of 2,975, 500 and 1,525 images with training, validation and test set. We train our network with training set, while evaluation is done using validation set. We resize images to 1024×512 for both training and testing as (Li, Yuan, and Vasconcelos 2019). We evaluate the class-level intersection over union (IoU) and mean IoU (mIoU) (Everingham et al. 2015).

Analysis

We first visualize each group through cross-domain grouping in Fig. 3. Clustered groups for various k showed that our networks clustered semantically consistent regions across the source and target domains through (1). Also, clustered regions along different k indicate that our network effectively divided regions into different group using (2).

We further compare our grouping network with k-means clustering algorithm (2012) and deep feature factorization (2018) which is not trainable methods. As shown in Fig. 4, our method can better capture the object boundaries and semantic consistent objects across source and target domains compared to other non-learnable methods. We further visualize each clustered group through cross-domain grouping with an evolving number of iteration. As the number of iterations increases, cross-domain grouping and group-level domain alignment share complementary information, which decomposes the data distribution and aligns domains for each grouped sub-spaces in a joint and boosting manner.

In Fig. 5, we show the t-SNE visualization (van der Maaten and Hinton 2008) of the output probability distribution of our method compares to the non-adapted method and baseline (Li, Yuan, and Vasconcelos 2019). The result shows that our method effectively aligns the distribution of

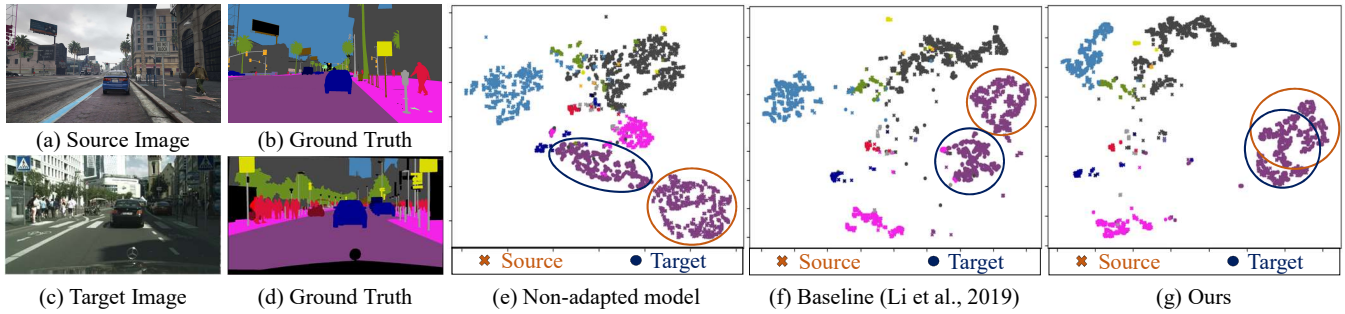


Figure 5: Visualization of output probability distribution of (a) source and (c) target image with GT classes as corresponding colors (b,d) via t-SNE using (e) non-adapted model, (f) baseline (Li, Yuan, and Vasconcelos 2019) and (g) ours. Our method effectively reduce domain discrepancy along with different domain, while others failed (represented using a circle).

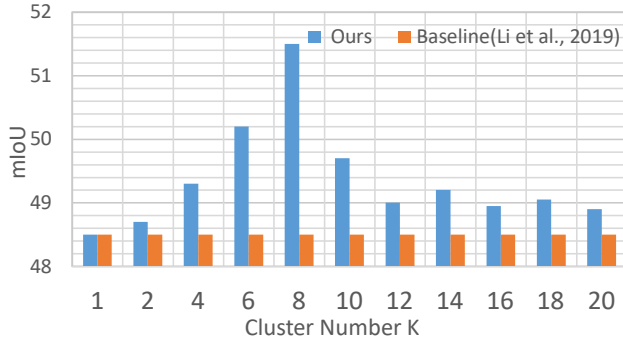


Figure 6: Ablation study for domain alignment with different number of clusters K on GTA5 \rightarrow Cityscapes.

the source and target domains, while others failed to reduce the domain discrepancy. Furthermore, we observe that our model successfully grouped minority categories (i.e. traffic signs in yellow) while others failed. It indicates that the loss (5) can solve class imbalance problem.

Ablation Study

Fig. 6 shows the result of ablation experiments with different number of groups K . Note that the results with $K = 1$ are equivalent to global domain adaptation as a baseline. The result shows that ours with various number of K consistently outperform the baseline, which shows the effectiveness of our group-level domain alignment. The performance has improved as K increased from 1, and after achieving the best performance at $K = 8$ the rest showed no significant difference. The lower performance with the larger number of K indicates that over-clustered samples can actually degrade performance as conventional category-level adaptation methods. Since the result with $K = 8$ has shown the best performance on both GTA5 \rightarrow Cityscapes and SYNTHIA \rightarrow Cityscapes, we set K as 8 for all experiments.

Table 1 shows the result of ablation experiments to validate the effects of proposed loss functions. It verifies the effectiveness of each loss function, including group-level domain adaptation, group-level semantic consistency, group-level orthogonality, and group-level class equivalence. The

Method	Loss Functions					mIOU
	L_{seg}	L_{cadv}	L_{co}	L_{orth}	L_{cl}	
Source only	✓					36.6
Ours	✓	✓				48.8
	✓	✓	✓			49.1
	✓	✓	✓	✓		50.8
	✓	✓	✓	✓	✓	51.5

Table 1: Ablation study for domain alignment with different loss functions on GTA5 \rightarrow Cityscapes.

full usage of our proposed loss functions yields the best results. We also find that adding group-level orthogonality leads to a large improvement in the performance, which demonstrates that we effectively divide the multi-modal complex distribution into K simple distributions for group-level domain alignment.

Comparison with State-of-the-art Methods

GTA5 \rightarrow Cityscapes. In the following, we evaluated our method on GTA5 \rightarrow Cityscapes in comparison to the state-of-the-art methods including without adaptation, global adaptation (Tsai et al. 2018), image-level adaptation (Wu et al. 2018; Chang et al. 2019; Li, Yuan, and Vasconcelos 2019) and category-level domain alignment (Luo et al. 2019; Du et al. 2019; Vu et al. 2019; Tsai et al. 2019; Huang et al. 2020; Wang et al. 2020). As shown in Table 2, our method outperforms all other models on the categories “car, truck, bus, motor, and bike” which share similar appearance. Specifically, we observe that our model achieves performance improvement on the categories “pole and traffic sign”. It demonstrates that our group-level class equivalence effectively solves the class imbalance problem.

SYNTHIA \rightarrow Cityscapes. We further compared our method to the state-of-the-art methods (Luo et al. 2019; Tsai et al. 2018; Du et al. 2019; Li, Yuan, and Vasconcelos 2019; Huang et al. 2020; Wang et al. 2020) on SYNTHIA \rightarrow Cityscapes, where 13 common classes between SYNTHIA (Ros et al. 2016) and Cityscapes (Cordts et al. 2016) datasets are evaluated. As shown in Table 3, our method outperforms conventional methods. We can observe that our

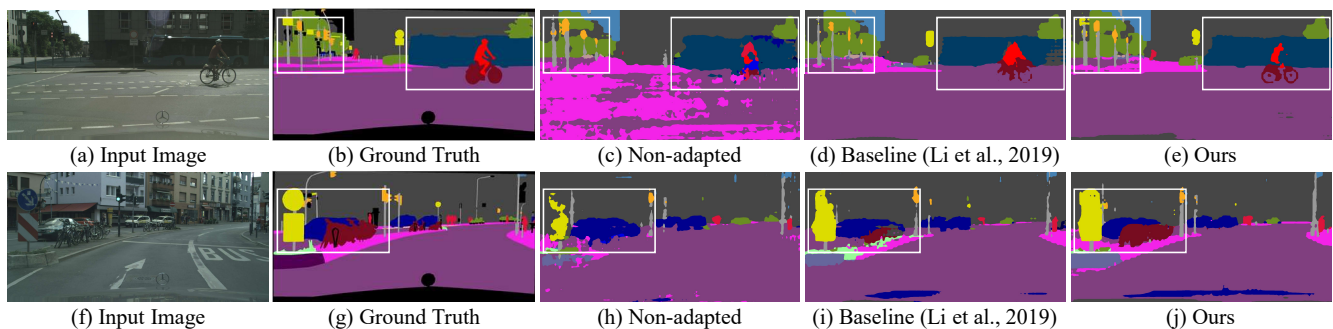


Figure 7: Qualitative results of domain adaptation on GTA5 \rightarrow Cityscapes. (From left to right) Input image, ground-truth, non-adapted result, baseline result and ours result.

GTA5 \rightarrow Cityscapes																				
	Road	SW	Build	Wall	Fence	Pole	TL	TS	Veg.	terrain	Sky	PR	Rider	Car	Truck	Bus	Train	Motor	Bike	mIoU
Without Adaptation	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
Tsai et al. (2018)	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
Wu et al. (2018)	85.0	30.8	81.3	25.8	21.2	22.2	25.4	26.6	83.4	36.7	76.2	58.9	24.9	80.7	29.5	42.9	2.5	26.9	11.6	41.7
Chang et al. (2019)	91.5	47.5	82.5	31.3	25.6	33.0	33.7	25.8	82.7	28.8	82.7	62.4	30.8	85.2	27.7	34.5	6.4	25.2	24.4	45.4
Li et al. (2019)	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
Luo et al. (2019)	87.0	27.1	79.6	27.3	23.3	28.3	35.5	24.2	83.6	27.4	74.2	58.6	28.0	76.2	33.1	36.7	6.7	31.9	31.4	43.2
Du et al. (2019)	90.3	38.9	81.7	24.8	22.9	30.5	37.0	21.2	84.8	38.8	76.9	58.8	30.7	85.7	30.6	38.1	5.9	28.3	36.9	45.4
Vu et al. (2019)	90.3	38.9	81.7	24.8	22.9	30.5	37.0	21.2	84.8	38.8	76.9	58.8	30.7	85.7	30.6	38.1	5.9	28.3	36.9	45.4
Tsai et al. (2019)	92.3	51.9	82.1	29.2	25.1	24.5	33.8	33.0	82.4	32.8	82.2	58.6	27.2	84.3	33.4	46.3	2.2	29.5	32.3	46.5
Huang et al. (2020)	92.4	55.3	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
Wang et al. (2020)	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
Ours	91.1	52.8	84.6	32.0	27.1	33.8	38.4	40.3	84.6	42.8	85.0	64.2	36.5	87.3	44.4	51.0	0.0	37.3	44.9	51.5

Table 2: Quantitative results of domain adaptation on GTA5 \rightarrow Cityscapes.

SYNTIA \rightarrow Cityscapes														
	Road	SW	Build	TL	TS	Veg.	Sky	PR	Rider	Car	Bus	Motor	Bike	mIoU
Tsai et al. (2018)	84.3	42.7	77.5	4.7	7.0	77.9	82.5	54.3	21.0	72.3	32.2	18.9	32.3	46.7
Li et al. (2019)	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
Luo et al. (2019)	82.5	24.0	79.4	16.5	12.7	79.2	82.8	58.3	18.0	79.3	25.3	17.6	25.9	46.3
Du et al. (2019)	84.6	41.7	80.8	11.5	14.7	80.8	85.3	57.5	21.6	82.0	36.0	19.3	34.5	50.0
Huang et al. (2020)	86.2	44.9	79.5	9.4	11.8	78.6	86.5	57.2	26.1	76.8	39.9	21.5	32.1	50.0
Wang et al. (2020)	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
Ours	90.7	49.5	84.5	33.6	38.9	84.6	84.6	59.8	33.3	80.8	51.5	37.6	45.9	54.1

Table 3: Quantitative results of domain adaptation on SYNTIA \rightarrow Cityscapes.

model achieve similar improvements in GTA5 \rightarrow Cityscapes scenario. Compared to baseline (Li, Yuan, and Vasconcelos 2019), our model achieves large improvement in the performance “traffic sign and traffic light”.

Conclusion

We have introduced cross-domain grouping and alignment for domain adaptive semantic segmentation. The key idea is to apply an intermediate grouping module such that multi-modal data distribution can be divided into several simple distributions. We then apply group-level domain alignment across source and target domains, where the grouping network and segmentation network can be trained in a joint and boosting manner using semantic consistency and or-

thogonality constraints. To solve the class imbalance problem, we have further introduced a group-level class equivalence constraint, resulting state-of-the-art performance on domain adaptive semantic segmentation. We believe our approach will facilitate further advances in unsupervised domain adaptation on various computer vision tasks.

Acknowledgements

This research was supported by R&D program for Advanced Integrated-intelligence for Identification (AIID) through the National Research Foundation of KOREA (NRF) funded by Ministry of Science and ICT (NRF-2018M3E3A1057289). (Corresponding author : Kwanghoon Sohn.)

References

- Arora, S.; Ge, R.; Liang, Y.; Ma, T.; and Zhang, Y. 2017. Generalization and equilibrium in generative adversarial nets (gans). In *International conference on machine learning*, 224–232.
- Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision*, 132–149.
- Chang, J.; Wang, L.; Meng, G.; Xiang, S.; and Pan, C. 2017. Deep adaptive image clustering. In *Proceedings of the IEEE International Conference on Computer Vision*, 5879–5887.
- Chang, W.-L.; Wang, H.-P.; Peng, W.-H.; and Chiu, W. 2019. All about structure: Adapting structural information across domains for boosting semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1900–1909.
- Chen, L. C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017a. 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.
- Chen, Y. H.; Chen, W. Y.; Chen, Y. T.; Tsai, B. C.; Wang, Y. C. F.; and Sun, M. 2017b. No more discrimination: Cross city adaptation of road scene segmenters. In *Proceedings of the IEEE International Conference on Computer Vision*, 1992–2001.
- Coates, A.; and Ng, A. Y. 2012. Learning feature representations with k-means. In *Neural networks: Tricks of the trade*, 561–580.
- Collins, E.; Achanta, R.; and Sussstrunk, S. 2018. Deep feature factorization for concept discovery. In *Proceedings of the European Conference on Computer Vision*, 336–352.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3213–3223.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Ding, C.; He, X.; and Simon, H. D. 2005. On the equivalence of nonnegative matrix factorization and spectral clustering. In *Proceedings of the SIAM International Conference on Data Mining*, 606–610.
- Du, L.; Tan, J.; Yang, H.; Feng, J.; Xue, X.; Zheng, Q.; Ye, X.; and Zhang, X. 2019. SSF-DAN: Separated semantic feature based domain adaptation network for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 982–991.
- Everingham, M.; Eslami, S. M. A.; Gool, L. J. V.; Williams, C. K. I.; Winn, J. M.; and Zisserman, A. 2015. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision* 111(1): 98–136.
- Ganin, Y.; and Lempitsky, V. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, 1180–1189.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2672–2680.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hoffman, J.; Tzeng, E.; Park, T.; Zhu, J. Y.; Isola, P.; Saenko, K.; Efros, A. A.; and Darrell, T. 2018. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, 1989–1998.
- Hoffman, J.; Wang, D.; Yu, F.; and Darrell, T. 2016. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. In *arXiv preprint arXiv:1612.02649*.
- Huang, J.; Lu, S.; Guan, D.; and Zhang, X. 2020. Contextual-relation consistent domain adaptation for semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, 705–722.
- Joulin, A.; Bach, F.; and Ponce, J. 2012. Multi-class cosegmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 542–549.
- karen Simonyan; and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. In *arXiv preprint arXiv:1409.1556*, –.
- Kingma, D. P.; and Ba, J. L. 2014. Adam: A method for stochastic optimization. In *arXiv preprint arXiv:1412.6980*, –.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 84–90.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11): 2278–2324.
- Li, G.; Kang, G.; Liu, W.; Wei, Y.; and Yang, Y. 2020. Content-consistent matching for domain adaptive semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, 440–456.
- Li, Y.; Yuan, L.; and Vasconcelos, N. 2019. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6936–6945.
- Long, J.; and Darrell, E. S. T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3431–3440.
- Long, M.; Cao, Z.; Wang, J.; and Jordan, M. I. 2018. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, 1640–1650.

- Luo, Y.; Zheng, L.; Guan, T.; Yu, J.; and Yang, Y. 2019. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2507–2516.
- Maas, A. L.; Hannun, A. Y.; and Ng, A. Y. 2013. Rectifier nonlinearities improve neural network acoustic models. In *International conference on machine learning*, 1–3.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch. –
- Richter, S. R.; Vineet, V.; Roth, S.; and Koltun, V. 2016. Playing for data: Ground truth from computer games. In *Proceedings of the European Conference on Computer Vision*, 102–118.
- Ros, G.; Sellart, L.; Materzynska, J.; Vazquez, D.; and Lopez, A. M. 2016. The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3234–3243.
- Shimodaira, H. 2000. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference* 90(2): 227–244.
- Tao, Z.; Liu, H.; Fu, H.; and Fu, Y. 2017. Image cosegmentation via saliency-guided constrained clustering with cosine similarity. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 4285–4291.
- Tsai, Y.-H.; Hung, W.-C.; Schuler, S.; Sohn, K.; Yang, M.-H.; and Chandraker, M. 2018. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7472–7481.
- Tsai, Y.-H.; Sohn, K.; Schuler, S.; and Chandraker, M. 2019. Domain adaptation for structured Output via discriminative Patch representations. In *Proceedings of the IEEE International Conference on Computer Vision*, 1456–1465.
- Tzeng, E.; Hoffman, J.; Saenko, K.; and Darrell, T. 2017. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7167–7176.
- van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9(Nov): 2579–2605.
- Vu, T.-H.; Jain, H.; Bucher, M.; Cord, M.; and Perez, P. 2019. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2517–2526.
- Wang, Z.; Yu, M.; Wei, Y.; Feris, R.; Xiong, J.; Hwu, W. M.; Huang, T. S.; and Shi, H. 2020. Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 12635–12644.
- Wu, Z.; Han, X.; Lin, Y.-L.; Uzunbas, M. G.; Goldstein, T.; Lim, S. N.; and Davis, L. S. 2018. Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In *Proceedings of the European Conference on Computer Vision*, 518–534.
- Yu, F.; and Koltun, V. 2015. Multi-scale context aggregation by dilated convolutions. In *arXiv preprint arXiv:1511.07122*, –.
- Zhao, H.; Shi, J.; Qi, X.; Wang, X.; and Jia, J. 2017. Pyramid scene parsing network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2881–2890.
- Zhao, X.; Liang, S.; and Wei, Y. 2018. Pseudo mask augmented object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4061–4070.
- Zou, Y.; Yu, Z.; Kumar, B. V.; and Wang, J. 2018. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European Conference on Computer Vision*, 289–305.