

# Coarse-to-Fine Domain Adaptive Semantic Segmentation with Photometric Alignment and Category-Center Regularization

Haoyu Ma<sup>1\*</sup> Xiangru Lin<sup>1\*</sup> Zifeng Wu<sup>2</sup> Yizhou Yu<sup>1†</sup>

<sup>1</sup>The University of Hong Kong <sup>2</sup>Deepwise AI Lab

mahaoyu@connect.hku.hk, xrlin2@cs.hku.hk, wuzifeng@deepwise.com, yizhouy@acm.org

## Abstract

*Unsupervised domain adaptation (UDA) in semantic segmentation is a fundamental yet promising task relieving the need for laborious annotation works. However, the domain shifts/discrepancies problem in this task compromise the final segmentation performance. Based on our observation, the main causes of the domain shifts are differences in imaging conditions, called image-level domain shifts, and differences in object category configurations called category-level domain shifts. In this paper, we propose a novel UDA pipeline that unifies image-level alignment and category-level feature distribution regularization in a coarse-to-fine manner. Specifically, on the coarse side, we propose a photometric alignment module that aligns an image in the source domain with a reference image from the target domain using a set of image-level operators; on the fine side, we propose a category-oriented triplet loss that imposes a soft constraint to regularize category centers in the source domain and a self-supervised consistency regularization method in the target domain. Experimental results show that our proposed pipeline improves the generalization capability of the final segmentation model and significantly outperforms all previous state-of-the-arts.*

## 1. Introduction

Semantic segmentation is a fundamental computer vision task that aims to assign a semantic category label to every pixel in an image. It has been widely used in many important downstream tasks such as autonomous driving [25, 5] and medical image analysis [4, 28, 10]. Recent state-of-the-art methods on semantic segmentation are primarily deep learning based [1, 15, 36] and require a large number of high quality annotated ground-truth data which are difficult to obtain especially in practical applications. Unsupervised domain adaptation semantic segmenta-

tion is an alternative method to solve the data scarcity problem where it generalizes models trained on the source domain composed of synthetic images and labels to perform well on the target domain composed of real world images only [21, 13, 30, 11, 20]. However, the problem is that semantic segmentation models trained merely on synthetic data exhibit poor performance on real world images due to the differences in multiple aspects (also called domain shifts/discrepancies), including exposure, contrast, lighting, object shape and surface textures, between the source domain and the target domain. Therefore, matching the distributions between the source and target domains to learn domain-invariant representations is crucial to solve the domain shifts.

Although the domain shifts could be caused by multiple factors, based on our observation, the primary causes can be summarized into two groups, namely image-level domain shifts and category-level domain shifts. For the image-level domain shifts, these refer to the differences in imaging conditions, such as lighting and settings in the camera imaging pipeline. Existing works on solving image-level domain shifts through image style transfer generally utilize deep models such as generative models and image-to-image translation models [17, 37] while another line of research focuses on using Fourier transformation [33]. These methods have proven that transferring image style of one domain to another domain can bring the two domains closer. However, the downside of these methods is that they either require to carry out a computationally expensive training process for the deep models or generate inferior style-transferred output images as shown in Figure 5.

Despite the fact that the domain gap can be minimized by global alignment methods such as the above, there is no guarantee that samples from different object categories in the target domain can be well separated. This is because some categories are naturally close to others in terms of body shape, pose and textures. To solve this problem, existing methods adopt category anchors computed on the source domain to guide the alignment between the two do-

\*These authors have equal contribution.

†Corresponding author

mains [34, 30], which can be regarded as a hard constraint on the category centers. The problem of this design is that it does not regularize the distance between different category features, and categories with similar feature distributions in the source domain also have similar distributions in the target domain, which results in erroneous classification results especially when no supervision information is available in the target domain. Our experimental results have demonstrated that imposing soft regularization methods on category distributions can improve the model’s capacity to adjust the relative magnitude of inter-category and intra-category feature distances.

According to the analysis above, performing alignment from either image-level perspective or category-level perspective alone will not solve the domain shifts reasonably. Therefore, we approach the problem from a different perspective and propose a novel and efficient pipeline that unifies image-level alignment and category-level feature distribution regularization in a coarse-to-fine manner. In general, on the coarse side, we propose a novel and efficient image-level alignment module to coarsely align the two domains; on the fine side, we introduce a new category-oriented triplet loss to softly regularize the category centers in the source domain and propose a self-supervised consistency regularization method in the target domain. By addressing both level of domain shifts simultaneously, we can significantly improve the performance of our proposed domain adaptation method.

**Coarse Alignment.** To solve the image-level domain shifts discussed above, we propose a global photometric alignment (GPA) module that aligns an image in the source domain with a reference image from the target domain using a set of image-level operators. Our method is superior to other generative methods and Fourier transformation based methods in two aspects: first, compared to the generative counterparts, our method requires no extra training process and produces stochastic image results; second, the quality of the translated image and the performance of our method is comparable to its generative counterpart and is superior to that of Fourier transformation based methods.

**Category-level Feature Distribution Regularization.** To address category-level domain shifts on the fine side, in addition to the common strategy of using pseudo labels for the target domain, we propose two novel regularization methods for the source and target domains respectively. First, considering the fact that there are annotated ground truth labels in the source domain, we propose a category-oriented triplet loss (CTL) that imposes a soft constraint to regularize category centers calculated using the source image pixel features, which actively enlarges the distances among category centers, making inter-category distances in a high-level feature space larger than intra-category distances by a predefined margin. Second, inspired

by the commonly used self-supervised learning methods: consistency regularization and pseudo-labeling, we propose a simple yet effective consistency regularizer for the target domain, called target domain consistency regularization (TCR), which constrains the prediction on an augmented target image to be consistent with the pseudo label of the corresponding non-augmented image, forcing the class labels of similar semantic contents to be consistent in the target domain.

In conclusion, this paper has the following contributions:

- We propose a novel coarse-to-fine domain adaptive semantic segmentation pipeline that seamlessly combines coarse image-level alignment with finer category-level feature distribution regularization.
- We introduce two novel and effective category-level regularization methods for the source and target domains respectively. The first one is called category-oriented triplet loss that regularizes category centers in the source domain while the second one performs target domain consistency regularization.
- Our method outperforms all previous methods, achieving new state-of-the-art performance on both GTA5→Cityscapes and SYNTHIA→Cityscapes benchmarks.

## 2. Related Work

Since our proposed domain adaptation pipeline is mostly related to photometric alignment based [35, 12, 33] and category-based domain adaptation methods [30, 34], we focus on these two types of work in this section.

**Photometric Alignment.** Previous works on domain adaptation [14, 26, 32, 31, 27, 16, 30, 8] have applied adversarial models, such as GAN [7, 17] and CycleGAN [37], to achieve photometric alignment results. Adversarial training makes a model capable of transferring image styles from one domain to another to significantly reduce the photometric differences between the two domains in the original image space [31, 30, 14]. Then a segmentation model trained on (style transferred) source domain images can be applied to target domain images [14, 30]. However, adversarial models are hard to train. Many researchers [9, 34, 30] have also shown that models based on adversarial training generally align distributions from different domains, but do not actually obtain mappings between features from different domains. Other types of photometric alignment methods for unsupervised semantic segmentation are rare. One method was proposed in [33] to align the source and target domains by simply replacing the low frequency component in a source domain image with its counterpart in the target domain reference image. However, such simple substitution of frequency components leaves unsatisfactory visual artifacts, and the performance of the model trained on the aligned samples relies heavily on a multi-band ensemble

ble. On the contrary, our method is different from previous method in that it has a light-weight photometric alignment strategy which does not require to carry out a computationally expensive training process and more importantly, produces comparable (superior) performance and image quality with respect to its generative (Fourier transformation based) counterpart.

**Category-Based Methods.** Category labels/predictions were introduced in [9, 16, 26, 29] to enforce global semantic constraints on the distribution of predicted labels. The proposed methods in [34] and [30] take one step further. They map penultimate target domain image features, that are used for generating pseudo labels in the output layer, to the corresponding features of the source domain image. However, in their work, category feature centroids [34] or instance features [30] in the source domain serve as anchors for category-based feature alignment, which does not explicitly enlarge the margins between the centers. This alignment strategy can be problematic because category anchors close in the source domain are likely hard to separate in the target domain as well. Our work differs from theirs in the following aspects: first, we propose a category-oriented triplet loss for the source domain that imposes a soft constraint to regularize category centers, actively making inter-category distances in a high-level feature space larger than intra-category distances by a specified margin; second, to further constrain category-level feature distributions in the target domain, we force the predictions on augmented target domain images to be consistent with the pseudo labels, generated by the segmentation model, of the corresponding non-augmented images, which is a self-supervision based consistency regularization method.

### 3. Method

#### 3.1. Coarse-to-Fine Pipeline

The key idea underlying our domain adaptation pipeline is intuitive: first, we exploit the photometric differences in the two domains and coarsely align the source domain images with the target domain images to minimize the domain shift; then, we regularize category-level feature distributions by setting constraints on inter-class center distances and intra-class feature variations.

**Step 0: Coarse Alignment.** Define  $\mathbf{M} = \{m_k\}_{k=1}^{N_s}$  as the source domain training set, where  $m_k$  is a source domain image and  $N_s$  is the number of images in the source domain training set. Similarly, the target domain training set is defined as  $\mathbf{N} = \{n_k\}_{k=1}^{N_t}$ . Our proposed GPA module converts a source domain image  $m$  in the training batch and a randomly selected target domain reference image  $n$  into Lab color space as  $(L_m, a_m, b_m)$  and  $(L_n, a_n, b_n)$ . Then the histogram mapping function  $f_{match}$  is applied to  $a_m$  and  $b_m$ , and gamma correction function  $f_{gamma}$  is applied to

$L_m$  to form  $(f_{gamma}(L_m), f_{match}(a_m), f_{match}(b_m))$ . The image is then converted to RGB space as aligned image  $m'$  to construct aligned source domain training set  $\mathbf{M}'$ . Then, a stochastic function  $\tau$  is applied to produce an augmented version of every image in  $\mathbf{M}'$ . A segmentation model  $T_0$  is trained based on all style-transferred images  $\tau(\mathbf{M}')$  with segmentation loss  $L_{seg}$ .

**Step 1: Category-level Feature Distribution Regularization.** In this step, we train a segmentation model  $T_1$  with  $\tau(\mathbf{M}')$  and  $\mathbf{N}$ . We apply the segmentation model  $T_0$  to all images in the target domain to produce a feature vector and a class probability vector at every pixel. The category corresponding to the largest value of the probability vector is defined to be the pseudo label at the pixel, and the largest probability value itself defines the confidence of the pseudo label. We further pre-define the pair of probability threshold  $P_h$  and percentage threshold  $p$  for all categories. The latter gives rise to a category specific probability threshold  $P_{s,c}$ , meaning  $p\%$  pixels in the category have confidence above  $P_{s,c}$ . Thus the final confidence threshold for category  $c$  is  $t_c = \min(P_h, P_{s,c})$ , and any pseudo labels in this category with a confidence higher than  $t_c$  are considered valid and added to the segmentation loss  $L_{seg}$ . The remaining pixels are left out during backpropagation. Then category center  $f_c$  for every category  $c$  are also calculated as the L2 normalized mean of all pixel features with category  $c$  as the ground truth label in the source domain. In addition to aligned training set  $\mathbf{M}'$  and cross-entropy loss  $L_{seg}$ , we impose a category-oriented triplet loss  $L_{triplet}$  on the segmentation model  $T_1$  in the source domain to enlarge inter-category distances, and a target domain consistency loss  $L_{consist}$  to regularize category-level feature distributions in the target domain. Then we finetune model  $T_0$   $U$  iterations to produce the model  $T_1$  by minimizing  $L_{seg} + L_{triplet} + L_{consist}$ .

**Step 2 to K: Iterative Self-Supervised Training.** Model  $T_1$  trained in Step 1 can be further improved with iterative steps similar to Step 1. Such an iterative approach is frequently called self-supervised training in the area of unsupervised domain adaptation for semantic segmentation [14, 34, 33, 30]. The same Step 1 is executed except that model  $T_{i-1}$  instead of  $T_0$  is used as the pre-trained model to generate pseudo-labels and category centers  $f_c$ . This process is repeated for  $K-1$  times. The overall pipeline of our proposed coarse-to-fine method is shown in Figure 1.

#### 3.2. Global Photometric Alignment

Since the global domain shift mostly affects low-level pixel attributes, which are irrelevant to pixel-wise category labels, we propose global photometric alignment (GPA) to align images from the source with images from the target domain. We observe that the spatial lightness distribution of an image can be very complicated under certain circum-

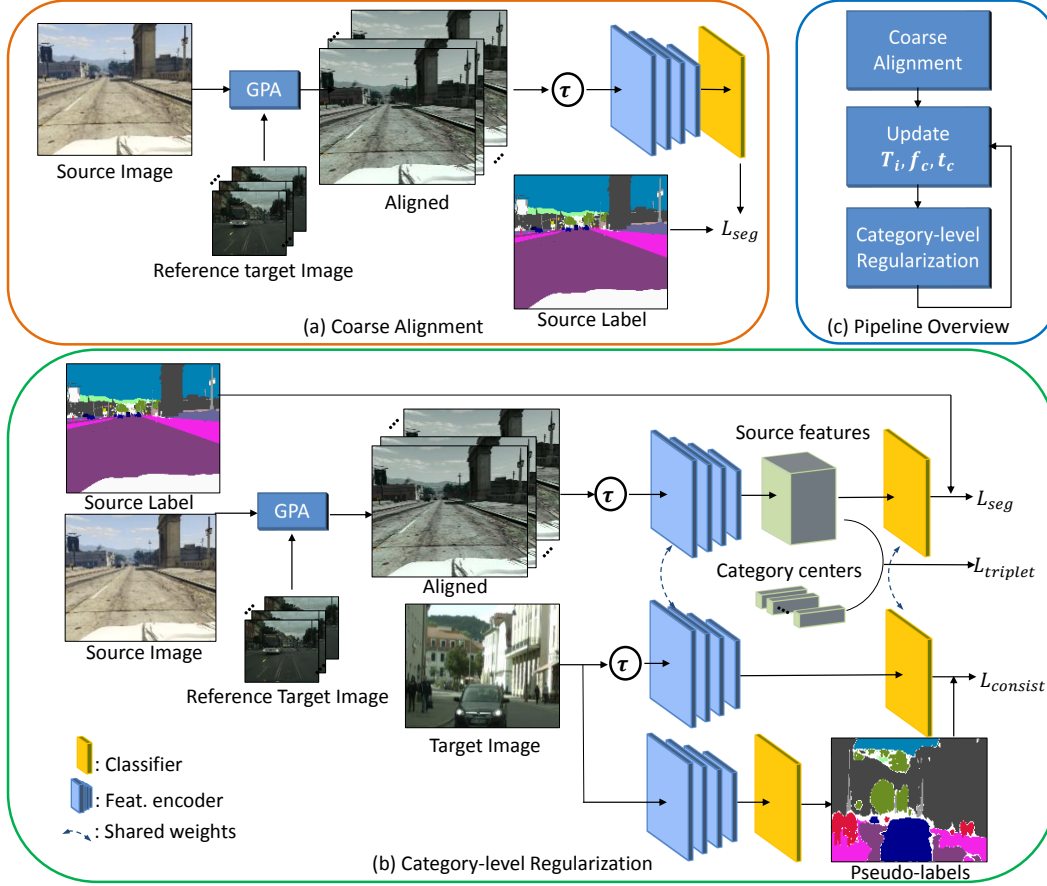


Figure 1. (a) First, the global photometric alignment (GPA) module is used to coarsely align the source and target domain images to train the initialized segmentation model  $T_0$ . (b) Then, we train the category-level feature distribution regularization step with the calculated category feature center  $f_c$  and pseudo-label threshold  $t_c$  for each category  $c$ . The category-oriented triplet loss is applied to the source domain and the consistency regularization is used in the target domain to jointly regularize the category-level feature distribution. (c) The overall pipeline is trained in an iterative self-supervised manner with  $T_i$ ,  $f_c$ , and  $t_c$  updated at each step  $i$ .

stances while the spatial color distribution of  $a$  and  $b$  have similar bell-shaped histograms. Therefore, we treat lightness and color differently and perform classic histogram matching [6] between the source domain image and the target domain reference image only on color channels  $a$  and  $b$  to avoid introducing artifacts commonly seen in histogram matching results.

**Lightness Gamma Correction.**  $L$  channel, on the other hand, is much more diversified among images. This is because light interacts with the 3D structure of a scene in a complicated manner. Simple histogram matching function gives rise to large areas of overexposure and fake structures. Thus, instead of strictly enforcing the mapping constraint prescribed by histogram matching for every histogram bin, we choose to only constrain the mean value of the lightness channel in the source domain image and make it equal to the mean value of the target domain reference image. Here, we choose the power-law function. The difference between our proposed method and the classic gamma correction is that

our function coefficients are automatically calculated with given source-target image pairs rather than user-defined. Specifically, we define  $f_{gamma}(L) = L^\gamma$ , where  $L$  is the normalized lightness value. Then the mean value constraint can be written as

$$\sum_L L p_s^m(f_{gamma}(L)) = \sum_L L p_s^m(L^\gamma) = \sum_L L p_t^n(L) \quad (1)$$

, where  $p_s^m$  is the lightness histogram of source image  $m$ , and  $p_t^n$  is the lightness histogram of target reference image  $n$ . This is a nonlinear equation and  $\gamma$  can be solved numerically.  $\gamma = 1$  when it is an identical transformation. In practice, to prevent  $\gamma$  from deviating too much away from 1, we introduce a regularization term into the following minimization,

$$\gamma^* = \arg \min_{\gamma} \left( \sum_L L p_s^m(L^\gamma) - \sum_L L p_t^n(L) \right)^2 + \beta(\gamma-1)^2, \quad (2)$$

which is a simple convex optimization problem with only one variable  $\gamma$ , and can be easily solved with few steps of gradient descent. The process of proposed GPA module is illustrated in Figure 2.




Figure 2. (a)Input source domain image and (b) a randomly chosen target domain image is aligned in (c) Lab channels to generate (d) aligned image.

### 3.3. Training Loss

The only training loss during coarse alignment step is the segmentation cross-entropy loss. The overall loss function we use during the category-level stage consists of three parts: the cross-entropy classification loss, a category-oriented triplet loss, and a target domain consistency regularization loss.

**Category-oriented Triplet Loss.** Even though the features learned with the GPA module are domain-invariant to some extent, the cross entropy losses used in previous training does not explicitly control the category-wise feature distribution. Therefore, the model learned with the GPA module using cross-entropy losses is coarsely aligned. Pixel features are distributed unevenly among different categories and some category centers are close to each other. To tackle this issue, we propose a category-oriented triplet loss that aims to further push the category-wise features closer to the corresponding category centers and further from other category centers. Let  $x_{i,j}$  be the pixel-wise features in the feature map of the second last layer, and  $y_{i,j}$  be the ground truth pixel-wise labels of a source domain image. The category center of category  $c$  is calculated as follows,

$$f_c = G\left(\frac{1}{N_c} \sum_s \sum_i \sum_j 1(y_{i,j} = c) x_{i,j}\right) \quad (3)$$

where  $N_c$  is the total number of pixels in category  $c$  and  $s$  is the source domain image index, and  $G$  is a L2 normalization function. Note that this L2 normalization is crucial to keep the category centers on the unit sphere to avoid scaling issue. The centers are updated after the training and this allows the centers become further and further from each other on the sphere surface.

Our category-oriented triplet loss is formulated as follows,

$$L_{triplet} = \frac{1}{N_s} \sum_s \sum_C \sum_i \sum_j \max(\|G(x_{i,j}) - f_{c=C}\| - \|G(x_{i,j}) - f_{c \neq C}\| + \alpha, 0), \quad (4)$$

where  $N$  is the total number of pixels in all images, and  $\alpha$  is a prescribed margin. The loss would be zero if every feature  $x_{i,j}$  is at least  $\alpha$  closer to its own category center than other category centers. Note that we only have reliable category labels for the source domain, thus we only apply the category-oriented triplet loss to the source domain images.




Figure 3. Our proposed category-oriented triplet loss exploits hard samples and further enlarge category margins.

The working principals of our proposed category-oriented triplet loss is illustrated in Figure 3. In cooperation with proposed photometric alignment and data augmentation in the source domain, our proposed triplet loss exploits hard samples in the coarsely aligned source domain and further improves the generalization capability of the trained model, which serves as complementary to cross-entropy loss.

#### Target Domain Consistency Regularization.

Category-oriented triplet loss is designed to regularize category-wise features in the source domain where the annotated ground truth labels are available. However, this is not the case in the target domain where no labels are provided. Consistency regularization is an important component of many recent state-of-the-art self-supervised learning algorithms, which utilizes unlabeled data by relying on the assumption that the model should output similar predictions when fed perturbed versions of the same image [23, 24]. Motivated by this, we propose a target domain consistency regularization method shown in Figure 1 to perform category-level feature distribution regularization in the target domain.

The idea of our proposed consistency regularization is simple: given a target domain image  $n_k$ , with the trained segmentation model  $T_{i-1}$ , we extract a pseudo label  $\hat{y}_j^k$  at every location  $j$  by feeding  $n_k$  to  $T_{i-1}$  followed by applying  $\arg \max(\cdot)$ ; and the corresponding pixel prediction is converted to a hard label vector  $\mathbb{1}_{[c=\hat{y}_j^k]}$ ; then, we apply the stochastic function  $\tau$  to  $n_k$  to obtain a perturbed version  $n'_k$ ; then, we feed  $n'_k$  to  $T_i$  to obtain a prediction  $p_j^{t_k}$  at every location  $j$  in the perturbed image; finally,  $p_j^{t_k}$  is forced to be consistent with  $\hat{y}_j^k$  by using a cross entropy loss function at pixel locations whose largest class probability is above the previously defined category-level confidence threshold  $t_c$ . By doing this, category-level feature distributions in the target domain are regularized under the supervision of valid

pseudo labels. The overall formula is defined as follows,

$$\begin{aligned}
 L_{cst} &= \sum_j \mathbf{1}(\max(T_{i-1}(n_k)|_j) \geq t_c) \text{CELoss}(\mathbf{1}_{[c=\hat{y}_j^k]}, p_j^{i_k}), \\
 \hat{y}_j^k &= \arg \max(T_{i-1}(n_k)|_j), \\
 p_j^{i_k} &= T_i(n'_k)|_j.
 \end{aligned} \tag{5}$$

It is important to use trained model  $T_{i-1}$  rather than model  $T_i$  to generate pseudo labels. This is because  $T_i$  is still being trained and unstable. Fluctuating pseudo labels generated by  $T_i$  would be catastrophic to the training process. Experimental results show this consistency regularization method is very effective even though the idea is simple.

## 4. Experiments

### 4.1. Datasets and Implementation Details

We follow the evaluation settings used in [34], and evaluate our proposed method with the source domain datasets **GTA5** [21] and **Synthia**[22], and the target domain dataset **Cityscapes** [3]. The **GTA5** dataset shares 19 common categories with the **Cityscapes** dataset and all the irrelevant categories are ignored during training; the **Synthia** dataset shares 16 common categories with the **Cityscapes** dataset. Some previous works only train and test on a 13-category subset of the **Synthia** dataset, or train two models on both subset and the whole set for better performance. Here we follow the practice in [18, 29] to train a model only on the whole set and test it on both settings.

According to Figure 1, we first use the photometrically aligned source domain images to train an initial segmentation model  $T_0$  in the coarse alignment step. Then, the model is trained in an iterative self-supervision manner with  $K = 6$  and  $U = 20k$ , and the total number of training iterations is  $140k$  which is comparable to all previous works [34, 30]. In our experiments,  $P_h = 0.9$  and  $p = 10$  for the pseudo-labels (as in [14]), and the regularization term  $\beta$  in (5) is 0.01. The margin  $\alpha$  is 0.2 for the triplet loss. We use the standard color-jittering as the stochastic function  $\tau(\cdot)$  in both source and target domains as in [29]. Following the same experimental settings in CAG [34], we adopt DeepLab V3+(Resnet101) [1]<sup>1</sup> as our segmentation model. Our proposed method has been implemented in PyTorch [19], and all experiments are conducted on 4 NVIDIA GeForce 2080Ti GPUs with 1 sample on each GPU. In the coarse alignment step, the stochastic gradient descent is used with momentum of 0.9 and weight decay of  $1e-4$ . The learning rate is initially set to  $5e-4$  and is decreased using the polynomial learning rate policy with power of 0.9. The setting for the following iterative finetuning steps are exactly the same except we halve the learning rate to  $2.5e-4$  to fine-tune previously trained models.

<sup>1</sup><https://github.com/RogerZhangzz/CAG.UDA/issues/6>

### 4.2. Comparison with State-of-the-Art Methods

In this section, we compare our method against all the existing state-of-the-art methods [18, 30, 34, 33, 29, 14], on both **GTA5**→**Cityscapes** and **Synthia**→**Cityscapes** tasks.

For the **GTA5**→**Cityscapes** task, according to Table 1, it is clear that our proposed method outperforms all previous methods, achieving a new state-of-the-art mIoU at 56.1% which is 5.9% higher than previous state-of-the-art methods [34]. In general, our method achieves the best performance in many important categories, including ‘road’, ‘sidewalk’, ‘building’, ‘light’, ‘sky’, ‘car’, ‘person’, ‘train’, ‘motor’, and ‘bike’. In particular, our model delivers a very good classification performance over ‘road’, ‘sidewalk’, ‘motor’ and ‘bike’ although some of these categories share very similar local appearances. This is because our category-oriented triplet loss focuses on the most confusing samples in different classes, and improves the generalization capability of the model. Moreover, the target consistency regularization in the target domain improves the classification accuracy of categories with a large intra-category variance, such as ‘building’ and ‘sky’.

The performance of the proposed method on **Synthia**→**Cityscapes** is shown in Table 2. The **Synthia** dataset has a larger domain shift caused by perspective and layout in addition to photometric differences in comparison to the **GTA5** dataset. But the overall performance of our model across all categories still surpasses the performance of other state-of-the-art methods, which demonstrates the effectiveness of our proposed techniques.

In comparison to CAG [34] using the same segmentation model, our proposed modules achieve a significant performance improvement, which is 5.9% in the **GTA5**→**Cityscapes** task and 3.7% in the **Synthia**→**Cityscapes** task. We further show some of the segmentation samples in Figure 4 to qualitatively demonstrate the superiority of our method. Please refer to the supplementary document for **more qualitative examples**.

### 4.3. Ablation Studies

**Component Analysis.** Most previous works [34, 18, 29, 30] require a segmentation model pre-trained on the original source domain training set only, and we call this model the source-only model. Although we do not use the source-only model during training, we train one to provide a baseline to demonstrate that the primary performance gain comes from our proposed modules and pipeline. As shown in Table 3, the performance of the source-only baseline using Deeplab v3+ is 37.6%, which is only slightly higher than that of the baseline using Deeplab v2 (36.6%) reported in [30, 26], and our proposed pipeline improves the baseline performance by 18.5%. Following the same settings in previous state-of-the-art methods [34, 18, 29], we further evaluate the impact of each proposed component on the performance of

	road	sidewalk	building	wall	fence	pole	light	sign	vege	terrace	sky	person	rider	car	truck	bus	train	motor	bike	mIoU
BDL [14]	91.0	44.7	84.2	34.6	27.6	30.2	36.0	36.0	85.0	<b>43.6</b>	83.0	58.6	31.6	83.3	35.3	49.7	3.3	28.8	35.6	48.5
IDA [18]	90.6	36.1	82.6	29.5	21.3	27.6	31.4	23.1	85.2	39.3	80.2	59.3	29.4	86.4	33.6	53.9	0.0	32.7	37.6	46.3
DTST [30]	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	<b>42.6</b>	48.5	1.9	30.4	39.0	49.2
FGGAN [29]	91.0	50.6	86.0	<b>43.4</b>	<b>29.8</b>	36.8	43.4	25.0	<b>86.8</b>	38.3	87.4	64.0	<b>38.0</b>	85.2	31.6	46.1	6.5	25.4	37.1	50.1
FDA [33]	<b>92.5</b>	53.3	82.3	26.5	27.6	36.4	40.5	38.8	82.2	39.8	78.0	62.6	34.4	84.9	34.1	<b>53.1</b>	16.8	27.7	46.4	50.4
CAG [34]	90.4	51.6	83.8	34.2	27.8	<b>38.4</b>	25.3	<b>48.4</b>	85.4	38.2	78.1	58.6	34.6	84.7	21.9	42.7	41.1	29.3	37.2	50.2
coarse align. (ours)	83.9	37.5	82.7	28.7	18.9	35.3	41.3	31.1	85.2	29.5	86.6	62.8	30.9	82.4	23.0	39.3	33.0	26.0	39.7	47.3
coarse-to-fine (ours)	<b>92.5</b>	<b>58.3</b>	<b>86.5</b>	27.4	28.8	38.1	<b>46.7</b>	42.5	85.4	38.4	<b>91.8</b>	<b>66.4</b>	37.0	<b>87.8</b>	40.7	52.4	<b>44.6</b>	<b>41.7</b>	<b>59.0</b>	<b>56.1</b>

Table 1. Performance comparison with state-of-the-art methods on the GTA5→Cityscapes task. Results after only coarse alignment and whole coarse-to-fine pipeline are both presented.




Figure 4. Qualitative analysis on GTA5→Cityscapes task. (a) Input images, (b) CAG [34], (c) Ours, (d) Labels

our model in the GTA5→Cityscapes task by removing one component at a time. According to our experimental results, the performance of the segmentation model has the most deterioration when the global photometric alignment module is removed. This is because in our coarse-to-fine pipeline, removing photometric alignment literally removes the first coarse alignment stage, and the resulting erroneous pseudo-labels are very detrimental to category-level feature distribution regularization. This also validates the necessity of a coarse alignment stage. Interestingly, although our target domain consistency regularization is simple, it has been proved to be very effective. This is because there are fewer training images in the target domain than the source domain, and filtering pseudo-labels with low confidence makes them even fewer. Our target domain consistency regularization increases the number of training samples in the target domain, therefore, giving rise to such a performance gain. The category-oriented triplet loss applied on the source domain also boosts the performance by 2.9% as it exploits hard samples in the source domain.

**Photometric Alignment.** There are currently other methods, which can achieve the goal of coarse alignment, such as the GAN-based method in [14, 2] and the frequency-based method in [33]. We substitute our proposed global photometric alignment with these two methods, and retrain our whole pipeline. The result is shown in Table 4. We also visualize some representative aligned

images produced with different methods in Figure 5. Our proposed GPA can generate the aligned image according to a randomly chosen target domain reference image, while the GAN-based model [14, 2] performs deterministically and generates aligned images with a similar style, only covering part of the actual target domain image span. This explains why our proposed model works even better than the pre-trained deep adversarial model. Although the frequency-based method proposed in [33] can generate style-transferred images randomly, the concatenation of frequencies usually introduces significant noises during training, which largely limits its final performance.

Based on our observation, gamma correction on all three channels does not have sufficient alignment capability, while histogram matching on all three channels results in image artifacts. We have run a comparison for the coarse alignment stage and the result (Table 4) shows our hybrid scheme performs the best.

**Pseudo-labels.** In our proposed method, we only apply the category-oriented triplet loss to source domain category labels but not pseudo-labels in the target domain. Although target domain images with pseudo-labels can be used as supplementary samples when the pseudo-labels are of high confidence, our proposed triplet loss aims to deal with hard samples, and pseudo-labels of hard samples in the target domain are not reliable. In order to verify this, we include pseudo-labels in our category-oriented triplet loss, and the

	road	sidewalk	building	wall	fence	pole	light	sign	vege	sky	person	rider	car	bus	motor	bike	mIoU	mIoU*
BDL [14]	<b>86.0</b>	<b>46.7</b>	80.3	-	-	-	14.1	11.6	79.2	81.3	54.1	27.9	73.7	42.2	25.7	45.3	-	51.4
IDA [18]	84.3	37.7	79.5	5.3	0.4	24.9	9.2	8.4	80.0	84.1	57.2	23.0	78.0	38.1	20.3	36.5	41.7	48.9
DTST [30]	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
FGGAN [29]	84.5	40.1	83.1	4.8	0.0	34.3	<b>20.1</b>	27.2	84.8	84.0	53.5	22.6	<b>85.4</b>	<b>43.7</b>	26.8	27.8	45.2	52.5
FDA [33]	79.3	35.0	73.2	-	-	-	19.9	24.0	61.7	82.6	61.4	31.1	83.9	40.8	<b>38.4</b>	51.1	-	52.5
CAG (13 classes) [34]	84.8	41.7	<b>85.5</b>	-	-	-	13.7	23.0	<b>86.5</b>	78.1	<b>66.3</b>	28.1	81.8	21.8	22.9	49.0	-	52.6
CAG (16 classes) [34]	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	44.5	-
coarse align. (ours)	64.0	25.7	73.9	9.6	0.8	33.3	12.3	25.9	81.6	85.5	62.4	26.2	80.6	30.9	26.8	23.8	41.5	47.7
coarse-to-fine (ours)	75.7	30.0	81.9	<b>11.5</b>	<b>2.5</b>	<b>35.3</b>	18.0	<b>32.7</b>	86.2	<b>90.1</b>	65.1	<b>33.2</b>	83.3	36.5	35.3	<b>54.3</b>	<b>48.2</b>	<b>55.5</b>

Table 2. Performance comparison with state-of-the-art methods on the Synthia→Cityscapes task (mIoU: 16-class; mIoU\*: 13-class).



Figure 5. Qualitative analysis on global photometric alignment. (a) Input images, (b) reference image, (c) BDL-GAN[14], (d) Fourier Adaptation[33], (e) Global photometric alignment.

	GPA	CTL	TCR	mIoU
Source only				37.6
w/o GPA		✓	✓	47.5
w/o CTL and TCR	✓			47.3
w/o CTL	✓		✓	53.2
w/o TCR	✓	✓		53.1
all	✓	✓	✓	56.1

Table 3. Ablation study of the proposed components on the GTA5→Cityscapes task. GPA: global photometric alignment, CTL: category-oriented triplet loss, TCR: target domain consistency regularization.

result is shown in Table 4.

Modules	Methods	mIoU
Image Align.	Frequency Align [33].	52.0
	BDL-GAN [14]	54.5
	Photometric Align.	56.1
GPA Scheme	Lab Gamma Correction	44.5
	Lab Histogram Match	43.3
	Hybrid	47.3
Pseudo-labels	Triplet loss with pseudo-labels	53.3
	Triplet loss w/o pseudo-labels	56.1

Table 4. Ablation studies of the image alignment plan, photometric alignment scheme, and using pseudo-labels for the category-oriented triplet loss on the GTA5→Cityscapes task.

## 5. Conclusions

In this paper, we propose a novel coarse-to-fine pipeline for domain adaptation semantic segmentation that smoothly integrates image-level alignment with category-level feature distribution regularization. In particular, we introduce a novel and efficient global photometric alignment module to coarsely align the source and target domains, and then, we propose a category-oriented triplet loss for the source domain and a target domain consistency regularization method to regularize the category-level feature distributions from a fine-grained category perspective. Experiments demonstrate that each of our proposed techniques improves the generalization capability of our model. And integrating them together results in a significant performance improvement in comparison to existing state-of-the-art unsupervised domain adapted semantic segmentation methods, demonstrating that solving image-level and category-level domain shifts simultaneously deserves more attention.

## Acknowledgments

This work was partially supported by National Key Research and Development Program of China (No.2020YFC2003902) and Hong Kong Research Grants Council through Research Impact Fund (Grant R-5001-18). H. Ma was supported by the Hong Kong PhD Fellowship.



## References

- [1] 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. [1](#), [6](#)
- [2] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1992–2001, 2017. [7](#)
- [3] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. [6](#)
- [4] Adrian V Dalca, Evan Yu, Polina Golland, Bruce Fischl, Mert R Sabuncu, and Juan Eugenio Iglesias. Unsupervised deep learning for bayesian brain mri segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 356–365. Springer, 2019. [1](#)
- [5] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 2020. [1](#)
- [6] Rafael C Gonzalez, Richard Eugene Woods, and Steven L Eddins. *Digital image processing using MATLAB*. Pearson Education India, 2004. [4](#)
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. [2](#)
- [8] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018. [2](#)
- [9] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. [2](#), [3](#)
- [10] Yuankai Huo, Zhoubing Xu, Yunxi Xiong, Katherine Aboud, Prasanna Parvathaneni, Shunxing Bao, Camilo Bermudez, Susan M Resnick, Laurie E Cutting, and Bennett A Landman. 3d whole brain segmentation using spatially localized atlas network tiles. *NeuroImage*, 194:105–119, 2019. [1](#)
- [11] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5001–5009, 2018. [1](#)
- [12] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. Tsit: A simple and versatile framework for image-to-image translation. *arXiv preprint arXiv:2007.12072*, 2020. [2](#)
- [13] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Chang-ick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6092–6101, 2019. [1](#)
- [14] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6936–6945, 2019. [2](#), [3](#), [6](#), [7](#), [8](#)
- [15] 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. [1](#)
- [16] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2507–2516, 2019. [2](#), [3](#)
- [17] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. [1](#), [2](#)
- [18] 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. [6](#), [7](#), [8](#)
- [19] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. [6](#)
- [20] Jiangtao Peng, Weiwei Sun, Li Ma, and Qian Du. Discriminative transfer joint matching for domain adaptation in hyperspectral image classification. *IEEE Geoscience and Remote Sensing Letters*, 16(6):972–976, 2019. [1](#)
- [21] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, pages 102–118. Springer, 2016. [1](#), [6](#)
- [22] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3234–3243, 2016. [6](#)
- [23] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in neural information processing systems*, pages 1163–1171, 2016. [5](#)
- [24] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020. [5](#)

- [25] Michael Trembl, José Arjona-Medina, Thomas Unterthiner, Rupesh Durgesh, Felix Friedmann, Peter Schuberth, Andreas Mayr, Martin Heusel, Markus Hofmarcher, Michael Widrich, et al. Speeding up semantic segmentation for autonomous driving. In *MLITS, NIPS Workshop*, volume 2, 2016. 1
- [26] 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. 2, 3, 6
- [27] 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. 2
- [28] Guotai Wang, Wenqi Li, Tom Vercauteren, and Sebastien Ourselin. Automatic brain tumor segmentation based on cascaded convolutional neural networks with uncertainty estimation. *Frontiers in computational neuroscience*, 13:56, 2019. 1
- [29] Haoran Wang, Tong Shen, Wei Zhang, Lingyu Duan, and Tao Mei. Classes matter: A fine-grained adversarial approach to cross-domain semantic segmentation. *arXiv preprint arXiv:2007.09222*, 2020. 3, 6, 7, 8
- [30] 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. 1, 2, 3, 6, 7, 8
- [31] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gokhan Uzunbas, Tom Goldstein, Ser Nam Lim, and Larry S Davis. Dcan: Dual channel-wise alignment networks for unsupervised scene adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 518–534, 2018. 2
- [32] Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 5423–5432, 2018. 2
- [33] 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. 1, 2, 3, 6, 7, 8
- [34] 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 435–445, 2019. 2, 3, 6, 7, 8
- [35] Yabin Zhang, Bin Deng, Kui Jia, and Lei Zhang. Label propagation with augmented anchors: A simple semi-supervised learning baseline for unsupervised domain adaptation. *arXiv preprint arXiv:2007.07695*, 2020. 2
- [36] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1529–1537, 2015. 1
- [37] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 1, 2