

Active Domain Adaptation with Multi-level Contrastive Units for Semantic Segmentation

Hao Zhang^{1,2*} and Ruimao Zhang^{**1}[0000-0001-9511-7532]

¹ Shenzhen Research Institute of Big Data, The Chinese University of Hong Kong(Shenzhen), China

² University of Illinois Urbana-Champaign
haoz19@illinois.edu
ruimao.zhang@ieee.org

Abstract. To further reduce the cost of semi-supervised domain adaptation (SSDA) labeling, a more effective way is to use active learning (AL) to annotate a selected subset with specific properties. However, domain adaptation tasks are always addressed in two interactive aspects: domain transfer and the enhancement of discrimination, which requires the selected data to be both uncertain under the model and diverse in feature space. Contrary to active learning in classification tasks, it is usually challenging to select pixels that contain both the above properties in segmentation tasks, leading to the complex design of pixel selection strategy. To address such an issue, we propose a novel Active Domain Adaptation scheme with Multi-level Contrastive Units (ADA-MCU) for semantic image segmentation. A simple pixel selection strategy followed with the construction of multi-level contrastive units is introduced to optimize the model for both domain adaptation and active supervised learning. In practice, MCUs are constructed from intra-image, cross-image, and cross-domain levels by using both labeled and unlabeled pixels. At each level, we define contrastive losses from center-to-center and pixel-to-pixel manners, with the aim of jointly aligning the category centers and reducing outliers near the decision boundaries. In addition, we also introduce a categories correlation matrix to implicitly describe the relationship between categories, which are used to adjust the weights of the losses for MCUs. Extensive experimental results on standard benchmarks show that the proposed method achieves competitive performance against state-of-the-art SSDA methods with 50% fewer labeled pixels and significantly outperforms state-of-the-art with a large margin by using the same level of annotation cost. Code will be in <https://github.com/haoz19/ADA-MCU>.

1 Introduction

Semantic segmentation is one of the most classic tasks in computer vision and image processing. The goal is to learn the semantic context in the image and

* Research done when Hao Zhang was a RA at SRIBD and CUHK,SZ

** Corresponding Author

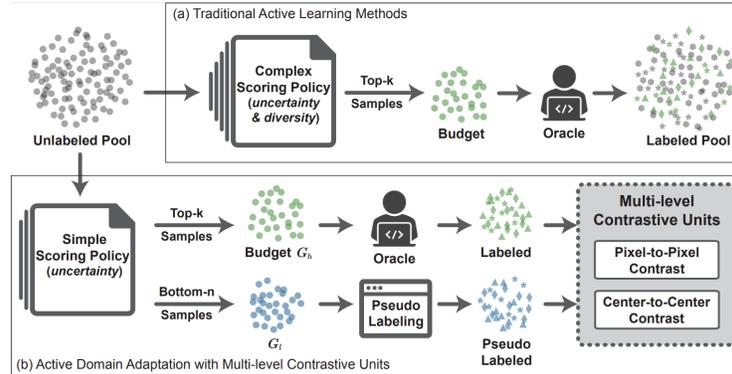


Fig. 1. traditional active learning methods via the complex sample selection policy for domain adaptation task and (b) the pipeline of our proposed Active Domain Adaptation with Multi-level Contrastive Units. G_l and G_h denote low and high uncertainty pixel groups determined by the uncertainty score.

automatically annotate each pixel a category label according to such learned information [1, 2]. In practice, such a task always requires excessively numerous annotations, limiting its scalability in real applications. One way to reduce annotation cost is to leverage a large amount of virtual data that is easy to obtain labels from game engines to extend training samples (*e.g.*, GTA5, SYNTHIA, Synscapes). However, the model trained merely with virtual data performs terribly on real-world data distribution because of the domain shifts. Therefore, many domain adaption methods are raised in recent years to bridge the gap between label-rich virtual data and label-scarce real-life data.

One way to address the domain adaptation problem is to train the model under the semi-supervised manner (SSDA), which jointly leverages fully labeled source domain data and a subset of labeled target domain data in the training phase. Compared to unsupervised domain adaptation (UDA), semi-supervised methods are able to significantly improve the segmentation accuracy through a small amount of annotated samples. Although the SSDA method makes a good balance between performance and annotation cost, there is still a big gap compared with the performance of the fully supervised approaches. In practice, how to bridge this gap by effectively labeling data is still an open issue.

In the literature, active learning (AL), which aims to select a subset of samples with specific properties, is proposed to annotate the training samples in a cost-effective way [19–21]. Although it has extensive research [23] on various areas by adopting active learning, it remains a challenge to deal with domain adaptation semantic segmentation since we need to address this task by dealing with two entangled issues from the pixel level, *i.e.*, **aligning domain distribution** and **improve the model’s discriminant ability**. In simple terms, we need to make the representative pixels (*i.e.*, the ones near the category center) from the same category but different domains be closer in the feature space, and to reduce uncertainty pixels (*i.e.*, the ones with low predictive confidence)

of each domain by pushing them to the corresponding centers from the decision boundary. However, it is especially challenging to select pixels that simultaneously meet the above two properties. Moreover, mining pixels being satisfied with the above properties from thousands of candidate pixels in an image through a complex scheme is also a time-consuming process.

To tackle such an issue, we propose a novel Active Domain Adaptation scheme with Multi-level Contrastive Units (ADA-MCU) for semantic segmentation. As shown in Fig. 1 (a) and Fig. 1 (b), different from active domain adaptation for image classification, which requires a complex scoring policy to labeled specific samples for model supervision, we use a simple selection policy in ADA-MCU to divide each image into two subsets of pixels, (*i.e.*, actively labeled ones with low confidence scores and unlabeled ones), then adopt pixels from these two subsets to construct contrastive pairs from multiple perspectives for model optimization. Specifically, at the **cross-domain** level, we enforce the distribution of the feature representations in the target domain being aligned to the source domain. While at the **intra-image** and **cross-image** level, we enforce the instance representations belonging to the same category to be closer in the feature space and make those are belonging to different categories to be far away in both source and target domain. Additionally, for each contrastive level, we do the alignment in two perspectives by using two kinds of loss, *i.e.*, center to center contrastive loss and pixel to pixel contrastive loss. The former enforces the centers of distribution to be aligned while the latter reduces the uncertainty of pixels by pulling them closer to the category center representations. In this way, two domains would be better aligned by employing the synergy of the above two losses.

In practice, since the misclassified pixels are highly relevant to its spatial layout, *e.g.*, the **sidewalk** is more likely to be misclassified as the **road** but be rarely misclassified as **sky**, we further introduce a dynamic categories correlation matrix (DCCM) to model the implicit relationship between each pair of categories. DCCM will be updated online during the training phase, aiming to adjust the weights of contrastive losses for the MCUs. Such a categories-aware contrastive loss could further improve the discriminative feature representation learning across domains. In this way, the domain transfer and discrimination enhancement are unified into one single framework.

The main contributions of this article can be summarized as follows,

1. We propose ADA-MCU, a novel active learning scheme, which uses a simple selection policy along with the construction of MCUs to optimize the model.
2. We introduce a simple yet effective scheme to construct Multi-level Contrastive Units (MCU) to regularize model training from multiple perspectives and propose a dynamic categories correlation matrix (DCCM) to describe the implicit relationship between categories, making more effective usage of the labeled and unlabeled pixels for model training.
3. As shown in Fig. 1 (c), extensive experiments demonstrate that our proposed method can achieve similar performance on two standard synthetic-to-real semantic segmentation benchmarks with less than 50% labeled data compared with current semi-supervised domain adaptation methods, and signif-

icantly outperforms state-of-the-art with a large margin by using the same level of annotation cost.

2 Related Work

2.1 Domain Adaptation for Semantic Segmentation

Domain adaptation task aims to explore how to transfer the knowledge that the model learned from one domain to another. In domain adaptation task, typically there are two domains, which have a certain domain shift but also share some common knowledge with each other. The domain that we train our model originally with is called source domain and the domain that we want to transfer our model to is called target domain. According to the usage of target domain annotations, we can divide domain adaptation semantic segmentation task into two main parts: 1) unsupervised domain adaptation one termed UDA and 2) semi-supervised domain adaptation termed SSDA. UDA is aimed at transferring the knowledge obtained from a labeled source domain to an unlabelled target domain. While SSDA aims at narrowing the gap with the fully supervised results by using a small set of labeled samples.

In recent years, some UDA methods for semantic segmentation [3–5] have been proposed via adversarial training which relies on a discriminator to measure the divergence between two domains’ distributions. Adversarial based methods aims at aligning the feature space of source domain and target domain by confusing the discriminator. Image translation has also been widely used in the UDA methods [6, 7]. As we know the most obvious gap between source domain and target domain is the color distribution. The aligning of color distribution is a very easy but efficient way to improve the performance. Recently some works [51] proposed that class-agnostic training paradigm gives model stronger generalization ability, which also pointed out a promising way of UDA. In addition, several methods [8, 9] also focus on self-training that generates pseudo labels for unlabelled data in the target domain.

For the SSDA, some methods [10–12] implemented feature alignment across domain from global and semantic level. For example, Chen *et al.* [11] proposed a framework based on dual-level domain mixing to address the differences in the amount of the labeled data between two domains. Huang *et al.* [12] aligned features by employing a few labeled target samples as anchors. However, none of those methods focus on sample selection, especially from the pixel-level perspective. We introduce pixel-level active learning to make more effective use of labeled pixels compared with SSDA.

Contrastive Learning. Contrastive learning (CL) aims at pushing positive sample pairs away from the negative ones in the representation space to learn better representations. For positive pair sampling, mainstream methods [13, 14] create different views of each sample using multiple perturbations. At the same time, negative pairs can be obtained by hard example mining strategies. [15, 16].

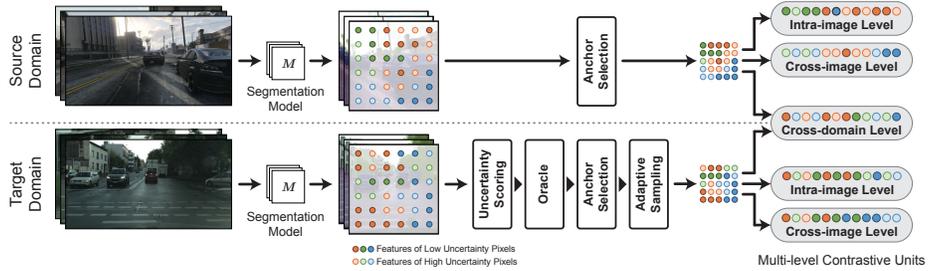


Fig. 2. Detailed illustration of construction of Multi-level Contrastive Units.

For the semantic segmentation task, CL could also be used to do intra-domain model pre-training [17]. Recently, [18] also propose a pixel-wise contrast scheme for semantic segmentation, making the representation of each category’s pixel more compact.

Active Learning. Active learning (AL) aims at obtaining high performance with low annotation costs by selecting the most informative samples. Usually, there are three major views to address AL tasks. 1) uncertainty-based methods [19, 20], 2) diversity-based methods [21, 22], 3) methods based on expected model change [23]. The former selected the samples with the highest uncertainty, while the latter selected the samples that could better represent the whole dataset. The last one aims to select the sample which can lead to more effect on the model. Compared to the existing works in active learning, our proposed method address AL in pixel-level and we construct MCUs for assistance in order to achieve great performance without a complicated selection strategy.

3 Methodology

In this section, we will describe our proposed method ADA-MCU in detail. Firstly, we will introduce a simple pixel-level sample selection policy, dividing each image into different pixel subsets. Secondly, we will present how to construct multi-level contrastive units (MCU) using labeled samples and unlabeled ones. At last, we will discuss how to apply contrastive loss with the dynamic category correlation to optimize the segmentation model.

3.1 Problem Setting and Notation

The goal of semantic segmentation domain adaptation is to transfer the model from source domain X_s to the target domain X_t . In our setting, we have a fully labeled source domain $\{(x_s^n, y_s^n)\}_{n=1}^{N_s}$, indicating the n -th source image x_s^n with the ground truth label map y_s^n , and the target domain $\{x_t^n\}_{n=1}^{N_t}$. Here N_s and N_t denote the number of source and target domain images, respectively. For the pixel-level active domain adaptation, the n -th target image x_t^n contains two

subsets in pixel-level, naming active annotated pixel $x_t^{n:(i,j)}$ with its corresponding ground-truth label $y_t^{n:(i,j)}$ and unlabeled pixel $x_t^{n:(\bar{i},\bar{j})}$, where (i, j) and (\bar{i}, \bar{j}) denote the labeled and unlabeled pixel positions in the target image. We use M_e and M_a to represent the number of expected labeled pixels number and already labeled pixels. And M_t is the total number of pixels in the target training dataset.

3.2 Active Pixel Annotation via Uncertainty Score

In practice, we first use source domain images to train the segmentation network. Once the pre-trained model \mathcal{F} is obtained, we use the output of \mathcal{F} on both source and target images to obtain the calculate their predictive uncertainty scores. The uncertainty score of pixel (i, j) in the n -th image can be calculated as follows,

$$S(x^{n:(i,j)}) = E(x^{n:(i,j)}) + \gamma D_{\text{KL}}(p^{n:(i,j)} || \hat{p}^{n:(i,j)}) \quad (1)$$

where the first part $E(\cdot)$ is the pixel-wise entropy, which is calculated by Eq. (2). The second part, $D_{\text{KL}}(\cdot)$ is the pixel-wise KL divergence between the predictions of the main segmentation head and the auxiliary head.³ γ is a hyperparameter to control the weights of two uncertainty indicators. $E(\cdot)$ and $D_{\text{KL}}(\cdot)$ can be calculated as follows,

$$E(x^{n:(i,j)}) = \frac{-1}{\log(C)} \sum_{c=1}^C p_c^{n:(i,j)} \log p_c^{n:(i,j)}, \quad (2)$$

$$D_{\text{KL}}(p || \hat{p}) = \sum_{c=1}^C p_c (\log p_c - \log \hat{p}_c), \quad (3)$$

where C denotes the total number of categories, By calculating the entropy of the pixel and the KL divergence of two head outputs, the degree of each pixel's uncertainty is obtained for further pixel annotation.

Using the scoring function mentioned above with the threshold π_{high}, π_{low} , we can divide both source and target image pixels into three groups termed *high*, *low*, and *medium* uncertainty groups respectively. Here the higher uncertainty indicates the lower model predictive confidence. By adjusting π_{high} , we can control the annotation rate of the target domain. We will discuss how to use pixels in different groups for contrastive unit construction in the next subsection.

Category Center Generation Strategy. After we divide both source and target pixels of each training batch into three groups, we use those pixels from the low uncertainty group with high predictive confidence to generate the category center. Intuitively, high confident samples always lie in the center of category

³ The auxiliary loss is proposed in [1] to improve the accuracy. We leverage the outputs from both auxiliary and main segmentation heads of DeepLab v2 to calculate KL divergence.

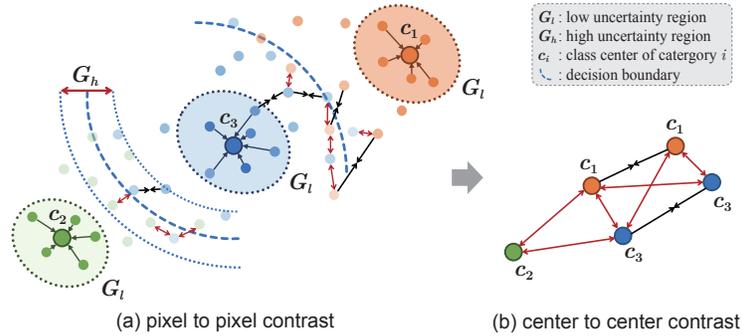


Fig. 3. Illustration of the construction of MCUs. We leverage both labeled data from G_h (when adding adaptive sampling, we also consider G_m to label) and unlabeled data from G_l to construct MCUs, and in each level of MCUs we use both pixel-to-pixel and center-to-center contrast to regularize the model in both uncertainty and diversity perspectives. Note that the representations can be from different pairs of images at different levels. High transparency means high uncertainty.

clusters, leading to a high density. According to the boundary assumption [36], the decision boundary should not cross the high-density region of a cluster. In other words, high confident pixels are always reliable and can be applied as representative of the cluster. Therefore, the aggregation of those pixels in the low uncertainty group for each domain can be considered as the representation of each category’s center.

Annotation via Adaptive Sampling When labeling the target domain pixels, we introduce adaptive sampling (AS) to maximize the advantages of active learning strategy. For each training batch, after dividing each target domain image into three groups: G_h^t , G_m^t and G_l^t , we randomly select the same number of pixels in those three groups. Then for G_h^t and G_m^t , we give ground-truth labels to those pixels, and we use the current predictions as pseudo labels to annotate those pixels selected by AS in G_l^t . Note that the size of each group changes for each training iteration. Thus the annotated pixels in G_h^t , G_m^t and G_l^t should be dynamically sampled, and we called it *annotation via adaptive sampling*. In practice, the size of G_m^t is always larger than G_h^t , therefore, in other words, we give a higher sampling rate to the group with high uncertainty for labeling.

3.3 Multi-level Contrastive Unit Construction

After obtaining the active labeled target pixels, we construct the multi-level contrastive units for domain adaptation and discrimination enhancement, which is partially inspired by the pixel-wise contrastive learning for semantic segmentation [25]. For each batch size that contains the same number of source and target images, we construct the contrastive units in three levels: (1) intra-image level, (2) cross-image level, and (3) cross-domain level.

Anchor Selection Policy. Before constructing the contrastive units for each level, we first require to select the anchor pixels for each contrastive unit to construct. In practice, rather than mining informative samples, we follow the work [25], determining whether it is a hard anchor through its uncertainty and whether the model predicts it correctly. In the source domain, we compare the predictive pixels’ labels and their ground truth and then select the ones which are incorrectly predicted as anchors. Similarly, in the target domain, we use active learning to annotate the pixels with high uncertainty. Then the incorrectly predicted pixels in the high uncertainty group⁴ are selected as the target anchors.

Contrastive Unit Construction. In each level, we firstly extract anchors from images using the above anchor selection policy. Then the pixels with the same label as the anchor are indicated as the positive samples, while the others with different labels are identified as the negative ones. After that, we construct the contrastive units in three levels, which are illustrated in Fig. 2. For the *intra-image level*, we extract positive pixels and negative pixels from the same image according to anchors. While there is a problem that the independence of images leads to the loss of information across the whole dataset. For example, the class `car` and class `train` may never appear in the same image through the entire domain, making these two categories lack inter-actions at the intra-image level. Therefore, we introduce *cross-image level* contrast to solving this problem. For the cross-images level, we first extract anchors from one image and extract positive pixels and negatives ones from another random image in the same domain. In this way, any two categories could be able to have interactions with each other along with the training iterations. According to the above strategy, the intra-image and the cross-image contrast units encourage the feature representation to be more discriminative in a specific domain. In contrast, we introduce *cross-domain level* contrastive units to align source and target distribution for the domain transfer. Different from cross-images level contrast, for cross-domain level, we extract anchors from one image in a certain domain and extract positive samples and negative ones from another image from the other domain. Through the synergy of three levels’ contrast units, we could achieve the discrimination enhancement and domain alignment in the model optimization process.

3.4 Pixel-level Active Domain Adaptation with MCUs

Segmentation Loss Function. In target domain, we use G_h^t , G_m^t , and G_l^t to denote the labeled pixels selected by adaptive sampling in high, medium, and low uncertainty groups, because there only labels for those samples selected by AL with AS in target domain. Thus the active training loss by using the labeled high uncertainty pixels can be defined as,

$$\mathcal{L}_{\text{seg}}^{G_h^t, G_m^t} = \sum_{n:(i,j) \in G_h^t, G_m^t} \mathcal{L}_{\text{CE}}(p^{n:(i,j)}, y^{n:(i,j)}) \quad (4)$$

⁴ Note that the pixel with high uncertainty does not necessarily to have the wrong predictive result.

where $p^{n:(i,j)}$ and $y^{n:(i,j)}$ are the prediction and ground truth of pixel in position (i, j) of n -th target image, and \mathcal{L}_{CE} denotes the cross-entropy loss function. On the contrary, pixels in G_l^t are with low uncertainty, which means more likely to be correctly predicted. Thus, we directly use their predictions as pseudo labels to calculate the loss,

$$\mathcal{L}_{\text{seg}}^{G_l^t} = \sum_{n:(\bar{i},\bar{j}) \in G_l^t} \mathcal{L}_{CE}(p^{n:(\bar{i},\bar{j})}, \hat{y}^{n:(\bar{i},\bar{j})}), \quad (5)$$

where $\hat{y}^{n:(\bar{i},\bar{j})} = \text{softmax}(p^{n:(\bar{i},\bar{j})})$. Thus the overall segmentation loss can be represented as follows,

$$\mathcal{L}_{\text{seg}} = \mathcal{L}_{\text{seg}}^{G_h^t, G_m^t} + \mathcal{L}_{\text{seg}}^{G_l^t} + \mathcal{L}_{\text{seg}}^{G_h^s, G_m^s, G_l^s}. \quad (6)$$

where G_h^s , G_m^s and G_l^s denote the pixel groups of the source domain. Note that since we have all annotations for images in source domain, we just calculate cross-entropy loss as their segmentation loss.

Dynamic Categories Correlation Matrix (DCCM). In practice, the context information and spatial layout are critical for segmentation accuracy. For instance, class `road` and `sidewalk` are misclassified more frequently, while `sky` and `road` are rarely to be misclassified. To better leverage such information, we introduce DCCM to describe the implicit relationship between any two categories. Denote $M_k^{(c_u, c_v)}$ as the number of pixels being misclassified from category c_u to c_v . Let M^{c_u} indicate the number of all pixels in c_u , then the error rate $R_k^{(c_u, c_v)}$ can be calculated as follow,

$$R_k^{(c_u, c_v)} = M_k^{(c_u, c_v)} / M^{c_u} \quad (7)$$

And at each iteration, we could dynamically update each element of the correlation matrix W by using exponential moving average as,

$$w_\tau^{(u,v)} = \beta w_{\tau-1}^{(u,v)} + (1 - \beta) R_k^{(c_u, c_v)}, \quad (8)$$

where $w_\tau^{(u,v)}$ denotes the correlation coefficient of u -th and v -th categories at the iteration τ . Then DCCM will further guide multi-level contrastive units by adjusting the weight of contrastive loss for each MCU.

Contrastive Loss Function. In each level of MCUs, we define loss function in two perspectives. On the one hand, we introduce pixel-to-pixel (**p2p**) contrastive loss based on the labeled pixels to reduce their uncertainty by pulling the same class pixel being close and pushing different class samples being apart. The loss function is defined based on InfoNCE [29], modified by using the weights in DCCM.

$$\mathcal{L}_{\text{con}}^{\text{p2p}} = \frac{1}{|D|} \sum -\log \mathcal{H}_p, \quad (9)$$

$$\mathcal{H}_p = \frac{\exp(w^{(u,v)} p \cdot p^+ / \lambda)}{\exp(w^{(u,v)} p \cdot p^+ / \lambda) + \sum_{p^- \in N_p} \exp(w^{(u,v)} p \cdot p^- / \lambda)}, \quad (10)$$

where \cdot denotes dot multiplication of two vectors with the scalar as the output. P_p and N_p denote pixel-wise embedding collections of the positive and negative samples.

On the other hand, we introduce class-to-class (c2c) contrastive loss by using the category centers introduced in the former section, which are most representative for each category. Then we define c2c contrastive loss in the same form with the p2p,

$$\mathcal{L}_{\text{con}}^{\text{c2c}} = \frac{1}{|P_c|} \sum_{c^+ \in P_c} \log \mathcal{H}_c, \quad (11)$$

where P_c and N_c denote class-wise embedding collections of the positive and negative samples, and \mathcal{H}_c has the same form as Eqn. 10 but uses class-wise anchors instead of pixel-wise ones. Using two kinds of contrastive losses as regular terms, we can encourage the features from the same category to be closer and from different categories to be further. And the DCCM can increase the weight between the categories which are more likely to be misclassified so as to guiding model optimization. In this way, the total loss of the proposed ADA-MCU can be presented as $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{seg}} + \mathcal{L}_{\text{con}}$, where \mathcal{L}_{con} is the sum of ($\mathcal{L}_{\text{con}}^{\text{p2p}} + \mathcal{L}_{\text{con}}^{\text{c2c}}$) from three levels.

Table 1. Experimental results on GTA5-to-Cityscapes compared with current SSDA, semi-supervised learning (SSL) methods. 19-class mIoU (%) scores are reported on Cityscapes validation set by using 1.7%, 3.4%, 6.8%, 16.8% labeled pixels from whole dataset.

Type	Method	Label Percentage (%)			
		1.7	3.4	6.8	16.8
Supervised	Image-wise	-	41.9	47.7	55.5
SSL	CutMix (bmvc20)	-	50.8	54.8	61.7
	DST-CBC (arxiv20)	-	48.7	54.1	60.6
	MME (cvpr19)	-	52.6	54.4	57.6
SSDA	MinEnt (cvpr19)	47.5	49.0	52.0	55.3
	AdvEnt (cvpr19)	44.9	46.9	50.2	55.4
	ASS (cvpr20)	50.1	54.2	56.0	60.2
	FDA (cvpr20)	53.1	54.1	56.2	59.2
	DDM (cvpr21)	-	61.2	60.5	64.3
	PCL (arxiv21)	54.2	55.2	57.0	60.4
	Ours	58.7	61.6	63.9	65.8

Table 2. Experimental results on SYNTHIA-to-Cityscapes compared with current SSDA and semi-supervised learning (SSL) methods. 13-class mIoU (%) scores are reported on Cityscapes validation set. Note that 1.7% pixels of 2795 images are at the same pixel number of 50 images.

Type	Method	Label Percentage (%)			
		1.7	3.4	6.8	16.8
Supervised	Image-wise	-	53.0	58.9	61.0
SSL	CutMix (bmvc20)	-	61.3	66.7	71.7
	DST-CBC (arxiv20)	-	59.7	64.3	68.9
	MME (cvpr19)	-	59.6	63.2	66.7
SSDA	MinEnt (cvpr19)	52.9	56.4	57.9	62.5
	AdvEnt (cvpr19)	51.4	55.2	59.6	62.6
	ASS (cvpr20)	60.7	62.1	64.8	69.8
	FDA (cvpr20)	58.5	62.0	64.4	66.8
	DDM (cvpr21)	-	68.4	69.8	71.7
	PCL (arxiv21)	61.2	63.4	65.2	70.3
	Ours	66.2	69.1	70.6	73.1

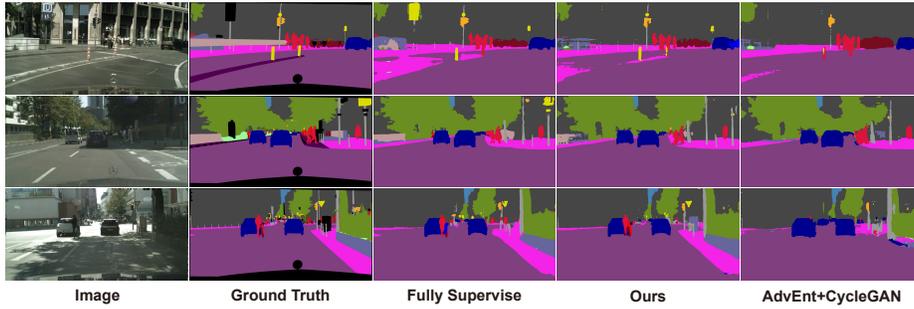


Fig. 4. Visualization of the segmentation results. Ours stands for our proposed method with 20% active annotations.

4 Experiment

4.1 Dataset, Setting and Implementation

We evaluate our proposed method by using the two standard large-scale segmentation benchmarks for domain adaptation, GTA5-to-Cityscapes, and SYNTHIA-to-Cityscapes. Following the previous method [32], we apply 19 classes domain adaptation for the former, and 13 classes for the latter. We conduct extensive experiments and report mean Intersection-over-Union (mIoU) compared with existing domain adaptation methods. All of the methods employ Deeplab v2 [2] as the basic model, which utilizes a pre-trained ResNet-101 [26] on ImageNet as backbone network. To measure the uncertainty, we calculate the KL divergence in Eqn. (3) by using multi-level outputs coming from both `conv4` and `conv5` feature maps. All experiments are run on a single Tesla V100 GPU with 32 GB of memory. All the models are trained by the Stochastic Gradient Descent (SGD) optimizer with an initial learning rate of 2.5×10^{-4} and decreasing with the polynomial annealing procedure with the power of 0.9.

Before the domain adaptation procedure, we use translated source domain data to train the model first. Then in each iteration, we randomly select 4 images, contains 2 from the source domain and 2 from the target domain. For **source domain images**, we feed them to the model and get the predictions and pixel-wise features. After that, we use predictions and labels to calculate cross-entropy loss, execute anchor selection and calculate intra-image and cross-image contrastive loss. While for **target domain images**, since we don't have any annotations at the beginning, we directly calculate the pixel-wise uncertainty score of the model output (after `softmax` function) and divide the batch of images into three groups according to the uncertainty, where γ is set to 0.5 in Eqn. (1). We ask the Oracle to label the high uncertainty group, and use the predictions as pseudo labels for pixels in low uncertainty group. Then, we can calculate cross-entropy loss and intra-image and cross-image contrastive loss similar to source domain. Furthermore, we also calculate cross-domain contrastive loss using both source domain and target domain data.

	AS	MCU- L_i	MCU- L_d	DCCM
AL(w/o AS)				
AL(w AS)	✓			
AL(w MCU $_i$)	✓	✓		
AL(w MCU $_d$)	✓	✓	✓	
Full Model	✓	✓	✓	✓

Table 3. Experiments setting of ablation study.

4.2 Comparison with State-Of-The-Art Methods

As presented in Table 1 and Table 2, we compare the proposed method with two SSL methods, *i.e.* CutMix [30], DST-CBC [31], and seven SSDA methods, *i.e.* MME [32], ASS [33], MinEnt [4], AdvEnt [4], FDA [35], PCL [34], DDM [11], in different percentage of annotation: 1.7%, 3.4%, 6.8%, 16.8%. As expected, compared with those methods, our purposed method has achieved a significant accuracy (mIoU) improvement.

From the Table 1, we can clearly see that our method can achieve comparable results with only about 50% of the annotations compared with other methods. Even compared with the state of art SSDA method DDM, we can still achieve similar performance using 30% fewer annotations. In addition, our proposed method can achieve similar performance with the fully supervised method using only 16.8% annotation of the whole target set. The visualization of the segmentation results of fully supervised method, UDA method [4] and our proposed method with 20% annotations are shown in Fig. 4. We can clearly see that our proposed scheme with 20% annotations can obviously improve the effect of some critical small areas, *e.g.* **sign**, **rider** and **person**.

5 Comparison with State-Of-The-Art ADA Methods

Recently, one article termed Multi-Anchor Active Domain Adaptation (MADA) [41], which is about active domain adaptation for semantic segmentation, has been accepted by ICCV2021 as the oral representation. The authors claim that it is the first study to adopt active learning to assist the domain adaptation regarding the semantic segmentation tasks, which adopts multiple anchors obtained via clustering-based method to characterize the feature distribution of the source-domain and multi-anchor soft-alignment loss to push the features of the target samples towards multiple anchors leading to better latent representation. Since the DeepLab v3+ [43] is adopted as the backbone network of this method, we have not listed in the main article considering fair comparisons with previous DeepLab v2 [2] based methods.

Different from our proposed *pixel-level* group partition regarding the uncertainty, such a method still adopts an *image-level* scheme to conduct active target sample selection against source anchors. As presented in Table 5, we compare our method with this active domain adaptation (ADA) method, MADA [41]. For

fairness, we follow the same setting that MADA uses. The DeepLab v3+ [43] is applied with the pre-trained ResNet-101 on ImageNet as the backbone network. We set $M_e = 5\%$ since MADA also selects 5% target-domain samples as active samples for their experiments. The other settings are the same as the former experiments described in the main article. The experiment shows that our proposed method outperform MADA by a large margin, *i.e.*, 1.6% mIoU, which demonstrates the proposed method could take little annotation workload but brings large performance gain compared with the recent SOTA method.

5.1 Ablation Study

Effectiveness of AL with Adaptive Sampling. Firstly, we want to investigate the effectiveness of our proposed active learning strategy. As shown in Table. 1, **UDA** is the performance of an unsupervised domain method AdvEnt with cycleGAN [4]. **RBA** denotes the performance of a pixel-level Region-based active learning method proposed in [37]. **AL(w/o AS)** indicates the degradation model which only uses the pixels from the high uncertainty group (*i.e.* pixels with manual labeling) and all of the pixels from the low uncertainty group (*i.e.* pixels with pseudo labels) to optimize the model. According to the results shown in Table. 1, we surprisingly find out that in our work, merely adding an active learning strategy may lead to performance degradation, *e.g.* 43.2% compared with 46.3%. One acceptable reason is that our proposed scheme not only uses labeled samples but also uses some of the unlabeled ones with pseudo labels to supervise the network, and the proportion of selected labeled pixels in different images varies greatly. For instance, if we set the annotation percentage for the whole dataset to be 10%, some of the images may get 2% or less labeled data (*e.g.* images with simple scenes) yet more than 50% pseudo label to supervise the model together. In such case, the loss calculated by those ground truth labels would be overwhelmed by the loss calculated by the pseudo labels, especially when the annotation budget is extremely limited. After we introduce adaptive sampling (AS) to our pixel-level active learning scheme, we get a satisfactory result (64.3% mIoU) shown by **AL(w AS)** in Table.1. Such result also outperforms **RBA** by 2.5%, showing the superiority of our proposed active selection strategy.

6 Conclusion

In this paper, we propose ADA-MCU, a novel active learning method, which uses a simple selection policy along with the construction of MCUs to optimize the segmentation model. As shown in Fig. 1, such a scheme abandons the complex sample selection policy in previous methods, leading to a more efficient active supervised training process. To the best of our knowledge, this work is the first study to conduct pixel-level annotation-based active domain adaptation for semantic image segmentation. The multi-level contrastive units (MCU), together with dynamic categories correlation matrix (DCCM), are carefully designed for

GTA5 to Cityscapes																				
	road	sidewalk	building	wall	fence	pole	light	sign	vege	terrace	sky	person	rider	car	truck	bus	train	motor	bike	mIoU
UDA (cvpr19)	92.4	52.7	83.8	32.4	24.1	30.7	33.2	25.7	83.7	35.1	85.1	58.2	27.4	85.5	37.1	41.9	2.1	25.2	22.6	46.3
RBA (wacv19)	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	61.8
AL(w/o AS)	90.4	34.9	82.3	30.0	23.4	27.4	31.9	21.9	84.0	38.5	77.8	58.4	25.0	84.8	26.9	34.9	1.5	27.6	19.8	43.2
AL(w AS)	96.4	75.7	86.8	40.3	42.0	47.4	46.1	65.4	87.9	44.0	84.3	68.6	44.9	91.5	66.7	72.6	53.9	41.9	64.7	64.3
AL(w MCU _i)	96.8	76.5	86.8	40.5	43.2	47.4	48.5	66.3	88.6	50.7	80.7	69.4	48.4	91.7	67.4	73.2	54.2	45.6	66.6	65.5
AL(w MCU _d)	97.1	77.4	87.8	42.1	43.9	48.1	47.4	65.3	87.4	55.1	82.9	72.1	49.1	91.2	70.4	73.1	55.3	45.7	66.3	66.1
Full Model	97.2	78.3	88.4	46.0	42.9	48.5	48.6	66.5	89.2	54.9	89.3	70.3	49.7	92.1	70.9	72.2	49.0	46.4	67.0	66.7

Table 4. Evaluation of different components of proposed method on GTA5-to-Cityscapes, with 20% labeled pixels except UDA.

GTA5 to Cityscapes (DeepLab v3+)																				
	road	sidewalk	building	wall	fence	pole	light	sign	vege	terrace	sky	person	rider	car	truck	bus	train	motor	bike	mIoU
MADA (iccv21)	95.1	69.8	88.5	43.3	48.7	45.7	53.3	59.2	89.1	46.7	91.5	73.9	50.1	91.2	60.6	56.9	48.4	51.6	68.7	64.9
Full Model	97.3	78.5	88.7	50.8	44.3	49.6	49.5	64.1	89.3	55.9	91.8	68.7	37.5	91.6	65.9	74.6	58.6	41.5	65.5	66.5

Table 5. Experimental results on GTA5-to-Cityscapes compared with current Active learning domain adaptation (ADA) method [41] with 5% annotations. The segmentation network used in the above experiment is DeepLab v3+ [43], which utilizes a pre-trained ResNet-101 on ImageNet as the backbone.

efficient active supervised model training, leading to many appealing benefits. (1) It enables the models to learn more compact feature representation for each category. (2) It could employ fewer annotations (16.8%) to achieve comparable performance with fully supervised method (65.3% mIoU). (3) It is effective for dealing with boundaries and small objects. Future work will combine the proposed scheme with more powerful architecture, *e.g.*, vision transformer, to explore more challenge tasks, such as panoptic segmentation.

Acknowledgements The work is supported in part by the Young Scientists Fund of the National Natural Science Foundation of China under grant No. 62106154, by Natural Science Foundation of Guangdong Province, China (General Program) under grant No.2022A1515011524, by CCF-Tencent Open Fund, by Shenzhen Science and Technology Program ZDSYS20211021111415025, and by the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese Univeristy of Hong Kong (Shenzhen).

References

1. H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.
2. L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2017.
3. J. Hoffman, D. Wang, F. Yu, and T. Darrell, "Fcns in the wild: Pixel-level adversarial and constraint-based adaptation," *arXiv preprint arXiv:1612.02649*, 2016.
4. T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez, "Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2517–2526.
5. J. Huang, D. Guan, S. Lu, and A. Xiao, "Mlan: Multi-level adversarial network for domain adaptive semantic segmentation," *arXiv preprint arXiv:2103.12991*, 2021.
6. J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "Cycada: Cycle-consistent adversarial domain adaptation," in *International conference on machine learning*. PMLR, 2018, pp. 1989–1998.
7. J. Huang, D. Guan, A. Xiao, and S. Lu, "Fsd: Frequency space domain randomization for domain generalization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6891–6902.
8. Y. Zou, Z. Yu, B. Kumar, and J. Wang, "Unsupervised domain adaptation for semantic segmentation via class-balanced self-training," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 289–305.
9. J. Huang, D. Guan, A. Xiao, and S. Lu, "Cross-view regularization for domain adaptive panoptic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 133–10 144.
10. Z. Wang, Y. Wei, R. Feris, J. Xiong, W.-M. Hwu, T. S. Huang, and H. Shi, "Alleviating semantic-level shift: A semi-supervised domain adaptation method for semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 936–937.
11. S. Chen, X. Jia, J. He, Y. Shi, and J. Liu, "Semi-supervised domain adaptation based on dual-level domain mixing for semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11 018–11 027.
12. J. Huang, D. Guan, A. Xiao, and S. Lu, "Semi-supervised domain adaptation via adaptive and progressive feature alignment," *arXiv preprint arXiv:2106.02845*, 2021.
13. Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3733–3742.
14. M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *arXiv preprint arXiv:2006.09882*, 2020.
15. P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," *arXiv preprint arXiv:2004.11362*, 2020.
16. J. Robinson, C.-Y. Chuang, S. Sra, and S. Jegelka, "Contrastive learning with hard negative samples," *arXiv preprint arXiv:2010.04592*, 2020.

17. K. Chaitanya, E. Erdil, N. Karani, and E. Konukoglu, “Contrastive learning of global and local features for medical image segmentation with limited annotations,” *arXiv preprint arXiv:2006.10511*, 2020.
18. W. Wang, T. Zhou, F. Yu, J. Dai, E. Konukoglu, and L. Van Gool, “Exploring cross-image pixel contrast for semantic segmentation,” *arXiv preprint arXiv:2101.11939*, 2021.
19. D. D. Lewis and J. Catlett, “Heterogeneous uncertainty sampling for supervised learning,” in *Machine learning proceedings 1994*. Elsevier, 1994, pp. 148–156.
20. T. Scheffer, C. Decomain, and S. Wrobel, in *International Symposium on Intelligent Data Analysis*. Springer, 2001, pp. 309–318.
21. S. D. Jain and K. Grauman, “Active image segmentation propagation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2864–2873.
22. S. C. Hoi, R. Jin, J. Zhu, and M. R. Lyu, “Semisupervised svm batch mode active learning with applications to image retrieval,” *ACM Transactions on Information Systems (TOIS)*, vol. 27, no. 3, pp. 1–29, 2009.
23. A. Vezhnevets, J. M. Buhmann, and V. Ferrari, “Active learning for semantic segmentation with expected change,” in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012, pp. 3162–3169.
24. Y. Siddiqui, J. Valentin, and M. Nießner, “Viewal: Active learning with viewpoint entropy for semantic segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9433–9443.
25. W. Wang, T. Zhou, F. Yu, J. Dai, E. Konukoglu, and L. Van Gool, “Exploring cross-image pixel contrast for semantic segmentation,” *arXiv preprint arXiv:2101.11939*, 2021.
26. K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
27. A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimpelshin, L. Antiga *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” *Advances in neural information processing systems*, vol. 32, pp. 8026–8037, 2019.
28. J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
29. A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.
30. G. French, T. Aila, S. Laine, M. Mackiewicz, and G. Finlayson, “Semi-supervised semantic segmentation needs strong, high-dimensional perturbations,” 2019.
31. Z. Feng, Q. Zhou, G. Cheng, X. Tan, J. Shi, and L. Ma, “Semi-supervised semantic segmentation via dynamic self-training and classbalanced curriculum,” *arXiv preprint arXiv:2004.08514*, vol. 1, no. 2, p. 5, 2020.
32. K. Saito, D. Kim, S. Sclaroff, T. Darrell, and K. Saenko, “Semi-supervised domain adaptation via minimax entropy,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 8050–8058.
33. Z. Wang, Y. Wei, R. Feris, J. Xiong, W.-M. Hwu, T. S. Huang, and H. Shi, “Alleviating semantic-level shift: A semi-supervised domain adaptation method for semantic segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 936–937.

34. W. Liu, D. Ferstl, S. Schuler, L. Zebedin, P. Fua, and C. Leistner, "Domain adaptation for semantic segmentation via patch-wise contrastive learning," *arXiv preprint arXiv:2104.11056*, 2021.
35. Y. Yang and S. Soatto, "Fda: Fourier domain adaptation for semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4085–4095.
36. O. Chapelle and A. Zien, "Semi-supervised classification by low density separation," in *International workshop on artificial intelligence and statistics*. PMLR, 2005, pp. 57–64.
37. T. Kasarla, G. Nagendar, G. M. Hegde, V. Balasubramanian, and C. Jawahar, "Region-based active learning for efficient labeling in semantic segmentation," in *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2019, pp. 1109–1117.
38. Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3733–3742.
39. K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9729–9738.
40. O. Chapelle and A. Zien, "Semi-supervised classification by low density separation," in *International workshop on artificial intelligence and statistics*. PMLR, 2005, pp. 57–64.
41. M. Ning, D. Lu, D. Wei, C. Bian, C. Yuan, S. Yu, K. Ma, and Y. Zheng, "Multi-anchor active domain adaptation for semantic segmentation," *arXiv preprint arXiv:2108.08012*, 2021.
42. L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," *arXiv preprint arXiv:1706.05587*, 2017.
43. L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 801–818.
44. Y.-H. Tsai, W.-C. Hung, S. Schuler, K. Sohn, M.-H. Yang, and M. Chandraker, "Learning to adapt structured output space for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7472–7481.
45. T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez, "Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2517–2526.
46. G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. 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*, 2016, pp. 3234–3243.
47. A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
48. R. D. Hjelm, A. Fedorov, S. Lavioie-Marchildon, K. Grewal, P. Bachman, A. Trischler, and Y. Bengio, "Learning deep representations by mutual information estimation and maximization," *arXiv preprint arXiv:1808.06670*, 2018.
49. X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," *arXiv preprint arXiv:2003.04297*, 2020.

50. T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
51. L. Qi, J. Kuen, Y. Wang, J. Gu, H. Zhao, Z. Lin, P. Torr, and J. Jia, “Open-world entity segmentation,” *arXiv preprint arXiv:2107.14228*, 2021.