

Conflict-Based Cross-View Consistency for Semi-Supervised Semantic Segmentation

Zicheng Wang^{1,3*} Zhen Zhao¹ Xiaoxia Xing³ Dong Xu² Xiangyu Kong³ Luping Zhou^{1†}
¹University of Sydney ²University of Hong Kong ³Samsung Research China-Beijing
 zwan4733@uni.sydney.edu.au {zhen.zhao, luping.zhou}@sydney.edu.au dongxu@cs.hku.hk
 {xx.xing, xiangyu.kong}@samsung.com

Abstract

Semi-supervised semantic segmentation (SSS) has recently gained increasing research interest as it can reduce the requirement for large-scale fully-annotated training data. The current methods often suffer from the confirmation bias from the pseudo-labelling process, which can be alleviated by the co-training framework. The current co-training-based SSS methods rely on hand-crafted perturbations to prevent the different sub-nets from collapsing into each other, but these artificial perturbations cannot lead to the optimal solution. In this work, we propose a new conflict-based cross-view consistency (CCVC) method based on a two-branch co-training framework which aims at enforcing the two sub-nets to learn informative features from irrelevant views. In particular, we first propose a new cross-view consistency (CVC) strategy that encourages the two sub-nets to learn distinct features from the same input by introducing a feature discrepancy loss, while these distinct features are expected to generate consistent prediction scores of the input. The CVC strategy helps to prevent the two sub-nets from stepping into the collapse. In addition, we further propose a conflict-based pseudo-labelling (CPL) method to guarantee the model will learn more useful information from conflicting predictions, which will lead to a stable training process. We validate our new CCVC approach on the SSS benchmark datasets where our method achieves new state-of-the-art performance. Our code is available at <https://github.com/xiaoyao3302/CCVC>.

1. Introduction

Among different vision tasks, semantic segmentation is a fundamental vision task that enables the network to understand the world [3, 11, 12, 32, 33]. In recent years, deep neu-

*This work was done during an internship at Samsung Research China-Beijing. This work is supported by Australian Research Council (ARC DP200103223).

†Corresponding authors.

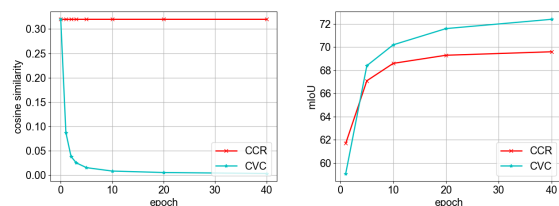


Figure 1. We compare the cosine similarity values between the features extracted by the two sub-nets of the traditional cross-consistency regularization (CCR) method and our CVC method. We also compare the prediction accuracies of the two methods, measured by mIoU. We show that our CVC method can prevent the two sub-nets from collapsing into each other and inferring the input from irrelevant views, while CCR cannot guarantee the inferred views are different. We show our new method can increase the perception of the model, which produces more reliable predictions. The experiments are implemented on the original Pascal VOC dataset, under the 1/4 split partition with ResNet-101 as the backbone of the encoder.

ral networks (DNNs) have shown great potential in semantic segmentation [18, 31, 57]. However, the success of DNNs is mainly due to the huge amount of annotated datasets. For the task of semantic segmentation, pixel-level annotations are often required, which means the annotators need to manually label up to hundreds of thousands of pixels per image. Therefore, it takes great effort to collect precisely labelled data for training DNNs [1, 27, 30].

Various semi-supervised learning (SSL) methods are proposed to tackle the problem, which aim at learning a network by using only a small set of pixel-wise precisely annotated data and a large set of unlabelled data for semantic segmentation [2, 34, 37, 53, 54, 58]. It is obvious that the information from the labelled data is very limited as the number of labelled data is far less than the number of unlabelled data. Therefore, it becomes a key issue to fully exploit the unlabelled data to assist the labelled data for the model training.

One intuitive way to tackle this issue is pseudo-

labelling [28,37,48]. However, SSL methods along this line may suffer from the so-called confirmation bias [48], which often leads to performance degradation due to the unstable training process. Recently, consistency regularization-based SSL methods show promising performance [35, 38, 41, 46]. However, most of them rely on producing the predictions of the weakly perturbed inputs to generate pseudo-labels, which are then used as the supervision to generate the predictions of the strongly perturbed inputs. Therefore, they still suffer from the confirmation bias issue.

On the other hand, co-training is a powerful framework for SSL as it enables different sub-nets to infer the same instance from different views and transfer the knowledge learnt from one view to another through pseudo-labelling. Particularly, co-training relies on multi-view reference to increase the perception of the model, thus enhancing the reliability of the generated pseudo-labels [40]. Various semi-supervised semantic segmentation (SSS) approaches are based on co-training [10, 39]. The key point is how to prevent different sub-nets from collapsing into each other such that we can make correct predictions based on the input from different views. However, the hand-crafted perturbations used in most SSS methods cannot guarantee heterogeneous features to be learned to effectively prevent sub-nets from stepping into a collapse.

Facing the above-mentioned issue, in this work, we come up with a new conflict-based cross-view consistency (CCVC) strategy for SSS, which makes sure the two sub-nets in our model can learn for different features separately so that reliable predictions could be learned from two irrelevant views for co-training, thus further enabling each sub-net to make reliable and meaningful predictions. In particular, we first raise a cross-view consistency (CVC) approach with a discrepancy loss to minimize the similarity between the feature extracted by the two sub-nets to encourage them to extract different features, which prevents the two sub-nets from collapsing into each other. Then we employ the cross pseudo-labelling to transfer the knowledge learnt from one sub-net to another to improve the perception of the network to correctly reason the same input from different views, thus producing more reliable predictions.

However, the discrepancy loss may introduce too strong a perturbation to the model that the feature extracted by the sub-nets may contain less meaningful information for the prediction, leading to inconsistent and unreliable predictions from the two sub-nets. This will incur the confirmation bias problem and thus harm the co-training of the sub-nets. To tackle this problem, we further propose a conflict-based pseudo-labelling (CPL) method, where we encourage the pseudo-labels generated by the conflicting predictions of each sub-net to have stronger supervision for the prediction of each other, to enforce the two sub-nets to make consistent predictions. Thereby, the useful features for the

prediction could be preserved as well as the reliability of the predictions. In this way, hopefully, the influence of the confirmation bias can be reduced and the training process can be more stable.

As shown in Fig. 1, we can see the similarity scores between the features extracted from the two sub-nets of the cross-consistency regularization (CCR) model remain at a high level, indicating the reasoning views of CCR are kind of relevant. In contrast, our CVC method ensures the reasoning views are sufficiently different and thus produces more reliable predictions.

It should be mentioned that our CCVC method is compatible with various existing data augmentation methods and it also benefits from an augmented training set with increased data diversity.

The contributions of our work are summarized as below:

- We introduce a cross-view consistency (CVC) strategy based on a co-training framework to make reliable predictions, where we propose a feature discrepancy loss to enable the two-branch network to learn how to reason the input differently but make consistent predictions.
- We further propose a new conflict-based pseudo-labelling (CPL) method based on our cross-view consistency strategy to enable the two sub-nets to learn more useful semantic information from conflicting predictions to produce reliable and consistent predictions, which leads to a more stable training process.
- Our method achieves the state-of-the-art performance on the commonly used benchmark datasets, PASCAL VOC 2012 [16] and Cityscapes [13].

2. Related work

2.1. Semantic segmentation

Semantic segmentation is a dense prediction vision task that aims at distinguishing the categories each pixel belonging to. FCN [36] is a pioneer work that proposed an encoder-decoder architecture with a fully convolutional network to perform pixel-wise semantic segmentation, which inspired tremendous works using a similar architecture to provide dense predictions, like the traditional convolutional neural network-based method including the DeepLab series [5–8], the HRNet [43], the PSPNet [52] and SegNeXt [21], etc. More recently, with the great success of the Transformer [4, 15, 42], various attempts have been proposed to utilize the great potential of the attention mechanism to capture the long-range contextual information for semantic segmentation like SegFormer [45], HRFormer [49], SETR [57] and SegViT [50], etc. However, the extraordinary performance of these methods re-

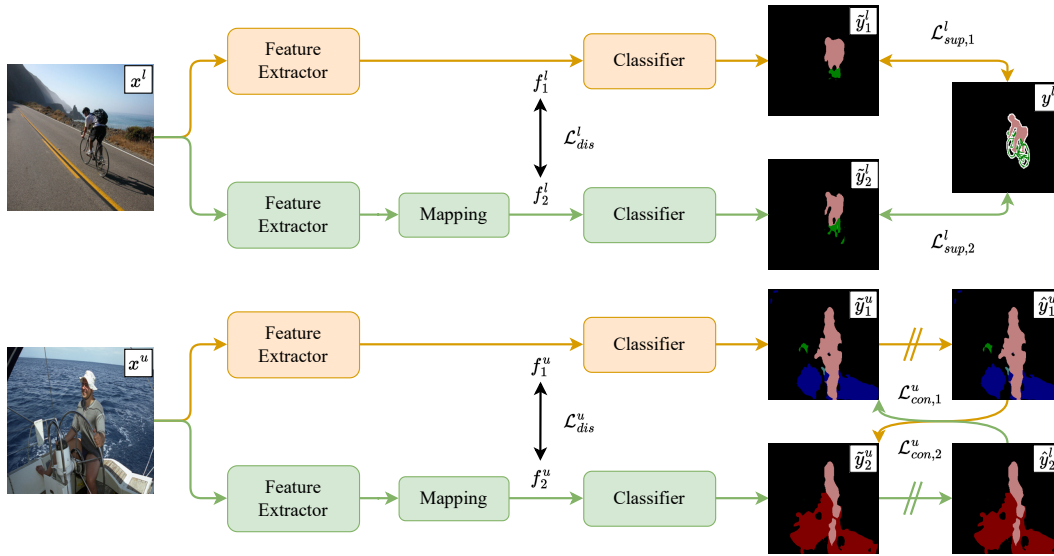


Figure 2. The network architecture of our cross-view consistency (CVC) method. We use a feature discrepancy loss to enforce the network to generate the same input from different views. On one hand, we use the supervised loss $\mathcal{L}_{sup,i}^l$ and the consistency loss $\mathcal{L}_{con,i}^u$ to perform cross-supervision. On the other hand, we use the discrepancy loss \mathcal{L}_{dis}^l to minimize the similarity between the features extracted by the feature extractors which thus enforcing the two sub-nets to learn different information. The subscript i denotes the i -th sub-net and the superscript α denotes the labelled data or the unlabelled data. The mark // represents the stop gradient operation. It should be mentioned that our CVC method is complementary to the traditional data augmentation methods.

lies heavily on full annotation supervision, and it is usually time-consuming to obtain the annotations.

2.2. Semi-supervised semantic segmentation

Semi-supervised learning (SSL) approaches were proposed to reduce the reliance of the model on large-scale annotated data. These methods aim at training a model using a small set of labelled data and a large set of unlabelled data. The key point of SSL methods is how to take full use of the large amount of unlabelled data [55, 56]. Current semi-supervised semantic segmentation methods can be roughly divided into two main categories, *i.e.*, self-training-based methods [10, 24, 28, 37, 44] and consistency regularization-based methods [29, 35, 38, 41, 59]. Most of the self-training-based methods select a set of predictions to generate pseudo-labels to fine-tune the model while most of the consistency regularization-based methods aim at using the network predictions of the weakly augmented inputs as the supervision for those predictions of the strongly augmented inputs. However, both of these two kinds of methods will suffer from a problem that the false positive predictions will introduce incorrect pseudo-labels and thus mislead the training, which is known as the confirmation bias.

2.3. Co-training

Co-training is a typical semi-supervised learning approach, which aims at learning two sub-nets to reason the

same instance from different views and then exchange the learnt information with each other [40, 47]. The two sub-nets can provide different and complementary information for each other, thus leading to a stable and accurate prediction and reducing the influence of the confirmation bias.

Various semi-supervised segmentation approaches have been proposed based on the co-training framework [10, 39], but the two sub-nets are easy to be collapsed. To prevent the issue, CCT [39] introduces feature-level perturbations to make sure the input of the several decoders is not the same to prevent the decoders from collapsing into each other. CPS [10] learns two sub-nets which are differently initialized, which ensures the two sub-nets behave differently.

However, it is difficult to guarantee the reasoning views of the different networks are different by using artificial perturbations. Compared with the above-mentioned methods, we propose to enable the networks to learn to be different but can still generate consistent predictions via our conflict-based cross-view consistency (CCVC) strategy, which utilizes a discrepancy loss to minimize the similarity of the features extracted by the two sub-nets to prevent the collapse and guarantee the reasoning views of the sub-nets are irrelevant.

3. Method

In this section, we will give a detailed explanation of our newly proposed conflict-based cross-view consistency

(CCVC) strategy. In Sec. 3.1 we will give a briefly definition of the semi-supervised semantic segmentation (SSS) task. Then, we will introduce our cross-view consistency (CVC) method in Sec. 3.2 and our conflict-based pseudo-labelling (CPL) method in Sec. 3.3.

3.1. Problem statement

In SSS tasks, we are given a set of fully pixel-wise annotated images $D_l = \{(x_i, y_i)\}_{i=1}^M$ and a set of unlabelled images $D_u = \{x_i\}_{i=1}^N$. M and N indicate the number of labelled images and unlabelled images. In most cases, we have $N \gg M$. The $x_i \in \mathbb{R}^{H \times W \times C}$ indicates the input image with a size of $H \times W$ and C channels, while $y_i \in \{0, 1\}^{H \times W \times Y}$ is the one-hot ground truth label for each pixel, where Y indicates the number of visual classes in total. We aim at training a model Ψ using the given input data D_l and D_u to generate semantic predictions.

3.2. Cross-view consistency

In this section, we illustrate our newly proposed cross-view consistency (CVC) method. We utilize a co-training-based two-branch network where the two sub-nets, *i.e.*, Ψ_1 and Ψ_2 , have a similar architecture but the parameters of the two sub-nets are not shared. The network architecture is shown in Fig. 2. Here, we divide each sub-net into a feature extractor $\Psi_{f,i}$ and a classifier $\Psi_{cls,i}$, where i equals 1 or 2, indicating the first sub-net and the second sub-net, respectively. Formally, we denote the feature extracted by the feature extractor $\Psi_{f,i}$ after L2 normalization as f_i^α and the prediction produced by the classifier $\Psi_{cls,i}$ as \tilde{y}_i^α , where $\alpha \in \{u, l\}$ represents for the labelled data stream or the unlabelled data stream, respectively. Recall that we aim at enabling the two sub-nets to reason the input from different views, so the feature extracted should be different. Therefore, we minimize the cosine similarity between the features f_i^α extracted by each feature extractor using a discrepancy loss \mathcal{L}_{dis}^α , which can be formulated as:

$$\mathcal{L}_{dis}^\alpha = 1 + \frac{f_1^\alpha \cdot f_2^\alpha}{\|f_1^\alpha\| \times \|f_2^\alpha\|} \quad (1)$$

Note that coefficient 1 is to ensure that the value of the discrepancy loss is always non-negative. We encourage the two sub-nets to output features with no co-relationship, thus enforcing the two sub-nets to learn to reason the input from two irrelevant views.

However, most SSS methods adopt a ResNet [23] pre-trained on ImageNet [14] as the backbone of the DeepLabv3+ and only fine-tune the backbone with a small learning rate, making it difficult to implement our feature discrepancy maximization operation. To tackle the issue, we follow a similar operation as BYOL [19] and SimSiam [9] to heterogeneity our network by mapping the features extracted by $\Psi_{f,2}$ to another feature space using a sim-

ple convolutional layer, *i.e.*, Ψ_{map} , with a non-linear layer. We denote the features extracted by $\Psi_{f,2}$ after mapping as \bar{f}_2^α and we re-write the discrepancy loss as:

$$\mathcal{L}_{dis}^\alpha = 1 + \frac{f_1^\alpha \cdot \bar{f}_2^\alpha}{\|f_1^\alpha\| \times \|\bar{f}_2^\alpha\|} \quad (2)$$

Note that we apply the discrepancy supervision on both the labelled data and the unlabelled data, so we calculate the total discrepancy loss as $\mathcal{L}_{dis} = 0.5 \times (\mathcal{L}_{dis}^l + \mathcal{L}_{dis}^u)$.

Note that we need to make sure the sub-nets make meaningful predictions. Therefore, for the labelled data, we use the ground truth label as supervision to train the two sub-nets to generate semantic meaningful predictions, and we formulate the supervised loss as follows:

$$\mathcal{L}_{sup,i}^l = \frac{1}{M} \sum_{m=1}^M \frac{1}{W \times H} \sum_{n=0}^{W \times H} \ell_{ce}(\tilde{y}_{mn,i}^l, y_{mn}^l) \quad (3)$$

Recall that the subscript i denotes the i -th sub-net and we use n to denote the n -th pixel in the m -th image, thereby \tilde{y}_{mn}^l and y_{mn}^l represents for the prediction or the ground truth label of the n -th pixel in the m -th labelled image, respectively. Note that we need to perform ground truth supervision on both of the two sub-nets, so we can calculate the supervised loss as $\mathcal{L}_{sup}^l = 0.5 \times (\mathcal{L}_{sup,1}^l + \mathcal{L}_{sup,2}^l)$.

For the unlabelled data, we adopt the pseudo-labelling approach to enable each sub-net to learn semantic information from the other one. Given a prediction $\tilde{y}_{mn,i}^u$, the pseudo label generated by it can be written as $\hat{y}_{mn,i}^u = \arg \max_c (\tilde{y}_{mnc,i}^u)$, where $\tilde{y}_{mnc,i}^u$ is the c -th dimension of the prediction score of $\tilde{y}_{mn,i}^u$ and $c = \{1, \dots, Y\}$ represents the index of the categories. We apply the cross-entropy loss to fine-tune the model, and the consistency loss for each branch can be formulated as below:

$$\mathcal{L}_{con,i}^u = \frac{1}{N} \sum_{m=1}^N \frac{1}{W \times H} \sum_{n=0}^{W \times H} \ell_{ce}(\tilde{y}_{mn,i}^u, \hat{y}_{mn,(3-i)}^u) \quad (4)$$

Recall that i equals 1 or 2, indicating the first or the second sub-net. The cross-consistency loss can be calculated as $\mathcal{L}_{con}^u = 0.5 \times (\mathcal{L}_{con,1}^u + \mathcal{L}_{con,2}^u)$.

To sum up, when learning the model, we jointly consider the supervised loss \mathcal{L}_{sup}^l , the consistency loss \mathcal{L}_{con}^u and the discrepancy loss \mathcal{L}_{dis} , the total loss can be written as follows:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{sup}^l + \lambda_2 \mathcal{L}_{con}^u + \lambda_3 \mathcal{L}_{dis} \quad (5)$$

where λ_1 , λ_2 and λ_3 are the trade-off parameters.

3.3. Conflict-based pseudo-labelling

With our cross-view consistency (CVC) method, the two sub-nets will learn from different views for semantic information. Nevertheless, the training might be unstable as the

feature discrepancy loss would introduce a too strong perturbation on the model. Thereby, it is hard to guarantee that the two sub-nets can learn useful semantic information from each other, which may further influence the reliability of the predictions.

To tackle the issue, we propose a conflict-based pseudo-labelling (CPL) method to enable the two sub-nets to learn more semantic information from the conflicting predictions to make consistent predictions, thereby guaranteeing that the two sub-nets can generate the same reliable predictions, and further stabilize the training. Here we use a binary value $\delta_{mn,i}^c$ to define whether the prediction is conflicting or not, where $\delta_{mn,i}^c$ equals 1 when $\hat{y}_{mn,1}^u \neq \hat{y}_{mn,2}^u$ and 0 otherwise. We aim at encouraging the model to learn more semantic information from these conflicting predictions. Therefore, when using these predictions to generate pseudo-labels for fine-tuning the model, we assign a higher weight ω_c to the cross-entropy loss supervised by these pseudo-labels.

However, the training may also be influenced by confirmation bias [48] during the training process as some of the pseudo-labels might be wrong. Therefore, following the previous methods [46, 51] that set a confidence threshold γ to determine whether the prediction is confident or not, we further divide the conflicting predictions into two categories, *i.e.*, the conflicting and confident (CC) predictions and the conflicting but unconfident (CU) predictions, and we only assign ω_c to those pseudo-labels generated by CC predictions. Here we use a binary value $\delta_{mn,i}^{cc}$ to define the CC predictions, where $\delta_{mn,i}^{cc}$ equals to 1 when $\hat{y}_{mn,1}^u \neq \hat{y}_{mn,2}^u$ and $\max_c(\tilde{y}_{mnc,i}^u) > \gamma$, otherwise $\delta_{mn,i}^{cc}$ equals to 0. Similarly, we can use $\delta_{mn,i}^e$ to denote the union of CU predictions and no-conflicting predictions where $\delta_{mn,i}^e = 1 - \delta_{mn,i}^{cc}$. It should be noticed that we still use the pseudo-labels generated with the CU predictions to fine-tune the model with a normal weight instead of directly discarding them, the main reason is that we argue that these CU predictions can also contain potential information about the inter-class relationship [44]. Therefore, we can re-write Eq. 4 as $\mathcal{L}_{con,i}^u = \omega_c \mathcal{L}_{con,i}^{u,cc} + \mathcal{L}_{con,i}^{u,e}$ where

$$\mathcal{L}_{con,i}^{u,cc} = \frac{1}{N} \sum_{m=1}^N \frac{1}{W \times H} \sum_{n=0}^{W \times H} \ell_{ce}(\tilde{y}_{mn,i}^u, \hat{y}_{mn,(3-i)}^u) \cdot \delta_{mn,i}^{cc} \quad (6)$$

and

$$\mathcal{L}_{con,i}^{u,e} = \frac{1}{N} \sum_{m=1}^N \frac{1}{W \times H} \sum_{n=0}^{W \times H} \ell_{ce}(\tilde{y}_{mn,i}^u, \hat{y}_{mn,(3-i)}^u) \cdot \delta_{mn,i}^e \quad (7)$$

Finally, we can re-calculate the total loss \mathcal{L} as calculated in Eq. 5 to train the model.

Our CCVC method can efficiently encourage the two sub-nets to reason the same input from different views and the knowledge transfer between the two sub-nets can in-

crease the perception of each sub-net, thus improving the reliability of the predictions.

It should be mentioned that in the inference stage, only one branch of the network is required to produce the prediction, and it also should be mentioned that our method is irrelevant to traditional data augmentation methods, which means we can directly adopt any data augmentation methods on the input data to increase the input diversity, the only thing to make sure is that the input of the two sub-nets should be the same.

4. Experiments

4.1. Datasets

Pascal VOC 2012 dataset [16] is a standard semi-supervised semantic segmentation (SSS) benchmark dataset, which consists of over 13,000 images from 21 classes. It contains 1,464 fully annotated images for training, 1,449 images for validation and 1,456 images for testing. Previous works use SBD [22] to render the labelled images and extend the number of labelled data to 10,582. The rendered labelled images are of low quality and some of them are accompanied by noise. Therefore, most of the previous works validate their SSS methods with sampled labelled images from the high-quality training images and rendered training images, respectively.

Cityscapes dataset [13] is another benchmark dataset for SSS, which focuses on urban scenarios and it consists of 2,975 annotated training images, 500 validation images and 1,525 testing images from 19 classes.

4.2. Implementation details

Following most of the previous works, we use DeepLabv3+ [8] as our segmentation model, which utilizes ResNet [23] pre-trained on ImageNet [14] as the backbone. Our mapping layer Ψ_{map} consists of a one-layer convolutional layer whose output dimension equals the input dimension, a BatchNorm layer [25], a ReLU function and a channel dropout operation with a dropout probability of 0.5. We use an SGD optimizer for our experiments with the initial learning rate set as 0.001 and 0.005 for Pascal VOC 2012 dataset and the Cityscapes dataset, respectively. We trained our model for 80 epochs and 250 epochs on the two datasets with a poly-learning rate scheduler, respectively (we only trained our model for 40 epochs on the original Pascal VOC 2012 dataset). The number of labelled data and unlabelled data are equal within a mini-batch and we set the batch size as 24 and 8 on the two datasets, respectively. We also adopt the weak data augmentation from PS-MT except for the crop size which we set as 512 and 712 for the two datasets, respectively. The weight ω_c of the consistency loss supervised by the pseudo-labels generated from the confident conflicting (CC) predictions is set

Table 1. Comparison with the state-of-the-art methods on the **PASCAL VOC 2012 dataset** under different partition protocols. Labelled images are from the original high-quality training set. The backbone is ResNet-101. The crop size of the input is set to 512. † indicates our model is only trained for 40 epochs while the other models are trained for 80 epochs.

Methods	1/16(92)	1/8 (183)	1/4 (366)	1/2 (732)	Full (1464)
Supervised Baseline	45.1	55.3	64.8	69.7	73.5
CutMix-Seg [17]	52.2	63.5	69.5	73.7	76.5
PseudoSeg [60]	57.6	65.5	69.1	72.4	73.2
PC ² Seg [58]	57.0	66.3	69.8	73.1	74.2
CPS [10]	64.1	67.4	71.7	75.9	-
ReCo [34]	64.8	72.0	73.1	74.7	-
ST++ [48]	65.2	71.0	74.6	77.3	79.1
U ² PL [44]	68.0	69.2	73.7	76.2	79.5
PS-MT [35]	65.8	69.6	76.6	78.4	80.0
Ours [†]	70.2	74.4	77.4	79.1	80.5

Table 2. Comparison with the state-of-the-art methods on the **PASCAL VOC 2012 dataset** under different partition protocols. Labelled images are sampled from the blended training set. The crop size of the input is set to 512. ‡ indicates our reproduced results.

Methods	ResNet-50			ResNet-101		
	1/16(662)	1/8 (1323)	1/4 (2646)	1/16(662)	1/8 (1323)	1/4 (2646)
Supervised Baseline	62.4	68.2	72.3	67.5	71.1	74.2
CutMix-Seg [17]	68.9	70.7	72.5	72.6	72.7	74.3
CPS [10]	72.0	73.7	74.9	74.5	76.4	77.7
CAC [29]	70.1	72.4	74.0	72.4	74.6	76.3
ELN [28]	-	73.2	74.6	-	75.1	76.6
ST++ [48]	72.6	74.4	75.4	74.5	76.3	76.6
U ² PL [‡] [44]	72.0	75.1	76.2	74.4	77.6	78.7
PS-MT [35]	72.8	75.7	76.4	75.5	78.2	78.7
Ours	74.5	76.1	76.4	77.2	78.4	79.0

as 2.0 for all of the experiments in this work, and we will give more discussion about the sensitivity of ω_c in the ablation study. In this work, we use the mean Intersection-over-Union (mIoU) as our evaluation metric. We set the hyper-parameters λ_1 , λ_2 and λ_3 as 5.0, 1.0, 2.0 on the original Pascal VOC 2012 dataset, 2.0, 2.0, 1.0 on the blended Pascal VOC 2012 dataset and 1.0, 1.0, 1.0 on the Cityscapes dataset. We set ω_c as 2.0 on all of the datasets and we set the pseudo-label threshold γ as 0.9 on the Pascal VOC 2012 datasets and 0.0 on the Cityscapes dataset.

4.3. Experimental results

We compare our CCVC method with recent semi-supervised semantic segmentation methods, including PseudoSeg [60], PC²Seg [58], CPS [10], ReCo [34], ST++ [48], U²PL [44] and PS-MT [35], etc. We also report the results of the re-implemented CutMix [17]. In addition, we also include the results of supervised methods that train the model with only labelled data for comparison (denoted as ‘‘Supervised Baseline’’). For all of the experiments, we follow CPS [10] and randomly split the datasets.

We first compare our methods with the others on the **original Pascal VOC 2012 dataset** and the results are reported in Table 1. Here we adopt the ResNet-101 as the

backbone of the encoder. We observe that our CCVC method achieves state-of-the-art (SOTA) results under all 5 partition protocols even though our model is only trained for 40 epochs while other models are trained for 80 epochs. It should be noticed that our method shows great power when the number of labelled data is small, *e.g.*, our method outperforms the current SOTA method by 2.2% and 2.4% when only 92 or 183 labelled data are available, respectively.

We further validate the effectiveness of our CCVC method on the **rendered Pascal VOC 2012 dataset** and the results are reported in Table 2. Here we report the results of using both the ResNet-50 and ResNet-101 as the backbone of the encoder respectively. We can see our CCVC method can also achieve SOTA results under all partition protocols when using different backbones, especially under the 1/16 partition protocol, our method surpasses the current SOTA method by 1.7% and 1.7% when using the ResNet-50 and ResNet-101 as the backbone, respectively, verifying the effectiveness of our method.

Finally, we test the performance of our CVC method on the challenging **Cityscapes dataset**. Due to the hardware memory limitation, we only report the results when using ResNet-50 as the backbone of the encoder. We can observe that even though the crop size of the input images is set

Table 3. Comparison with the state-of-the-art methods on the **Cityscapes dataset** under different partition protocols. The backbone is ResNet-50 and the crop size of the input is set to 712. * indicates U²PL reproduced results.

Methods	1/16 (186)	1/8 (372)	1/4 (744)
Supervised Baseline	63.3	65.8	68.4
CCT [39]	66.4	72.5	75.7
GCT [26]	65.8	71.3	75.3
CPS* [10]	69.8	74.3	74.6
ELN [28]	-	70.3	73.5
ST++ [48]	-	72.7	73.8
U ² PL [44]	69.0	73.0	76.3
USRN [20]	71.2	75.0	-
PS-MT [35]	-	75.8	76.9
Ours	74.9	76.4	77.3

as 712 and the training epoch is set as 250 in our work, our method still achieves the new SOTA performance, especially when there are only 186 labelled data available, our method surpasses the current SOTA method ST++ by 3.7%. In addition, our method surpasses U²PL, which sets the crop size as 769, and PS-MT, which trains the model for 450 epochs, verifying the effectiveness of our method.

It can be inferred from the tables that our method can achieve great performance especially when the number of labelled data is small, indicating that our method can take better use of the unlabelled data.

4.4. Ablation study

In this section, we analyze the effectiveness of the detailed module designs of our CCVC approach. Here we conduct all of the ablation experiments with a ResNet-101 as the backbone of the DeepLabv3+ on the original Pascal VOC 2012 dataset under the partition of 1/4.

Effectiveness of Components. Recall that our CCVC method includes a CVC module, a CPL module and data augmentation (Aug.). Note that there are three losses, *i.e.*, the supervised loss \mathcal{L}_{sup} , the consistency loss \mathcal{L}_{con} and the discrepancy loss \mathcal{L}_{dis} , used in our CVC method together with one extra mapping module Ψ_{map} . We now investigate the individual contributions of these losses and modules in CCVC. The analysis results are reported in Table 4.

We can observe that if we only apply the consistency loss, the cross-consistency regularization (CCR) method will bring a performance improvement of over 4%, but the reasoning views of the two sub-nets are kind of correlated, leaving a huge space for improvement. When applying our discrepancy loss, we can see that there is a 1.7% improvement of mIoU, indicating that our feature discrepancy maximization operation can ensure the two sub-nets reason the input from two orthogonal views, thus improving the perception of the model. In addition, when applying the mapping module, the reasoning views of the two sub-nets can be

Table 4. Ablation study on the effectiveness of different components in our CCVC method, including the supervised loss \mathcal{L}_{sup} , the consistency loss \mathcal{L}_{con} , the discrepancy loss \mathcal{L}_{dis} , the mapping operation Ψ_{map} , the conflict-based pseudo-labelling (CPL) and the data augmentation (Aug.) strategy. Note that we only train the model for 40 epochs.

\mathcal{L}_{sup}	\mathcal{L}_{con}	\mathcal{L}_{dis}	Ψ_{map}	CPL	Aug.	mIoU
✓						64.8
✓	✓					69.6
✓	✓	✓				71.3
✓	✓		✓			71.0
✓	✓	✓	✓			72.4
✓	✓	✓	✓	✓		74.0
✓	✓	✓	✓	✓	✓	77.4

further separated, enhancing the network to produce more reliable predictions. Therefore, the improvement will be further enhanced by 1.1%. It also should be noticed that when introducing the mapping module to the original co-training framework, the hand-crafted network perturbation can also reduce the collapse to a degree, leading to a performance improvement of 1.4%. In addition, when applying the CPL module to our CVC method, there is a giant improvement of the performance by 1.6%, verifying our hypothesis that the features learnt by the two sub-nets might be pushed far away to contain useful semantic information. Our CPL method can enable the sub-nets to learn more information from conflicting predictions, thus guaranteeing the predictions are reliable. When we use some simple strong data augmentations as listed in ST++ [48] to increase the diversity of the input data, our method can be further enhanced with an improvement of mIoU by 3.4%, which surpasses the SOTA methods by a large margin, verifying the effectiveness of our method.

We further verify that our CVC method can reduce the influence of confirmation bias. We compared our CVC method with the supervised baseline and cross-consistency regularization (CCR) method. Here we count the proportion of reliable predictions of each method as well as the corre-

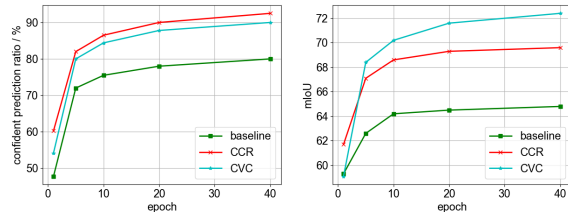


Figure 3. The training process of the supervised baseline, the cross-consistency regularization (CCR) method and our CVC method. It can be seen that our methods will not output many highly confident predictions, but the accuracy of the predictions is very high, indicating that our method can efficiently reduce the influence of the confirmation bias.

Table 5. Ablation study on the effectiveness of our conflict-based pseudo-labelling (CPL) method. We vary the weight, *i.e.*, ω_c , of the confident conflicting (CC) predictions to verify the model will learn more semantic information from conflicting predictions. Note that we only train the model for 40 epochs.

ω_c	1.0	1.2	1.5	1.8	2.0	2.5
mIoU	72.4	72.9	73.5	73.7	74.0	73.8

sponding mIoU during the training process. The threshold of the reliable prediction is set as 0.9. The results are listed in Fig. 3. We can observe that the CCR method can generate more confident predictions than our CVC method, while the performance of our CVC method is better than the CCR method, indicating that the CCR method will generate more confident but incorrect predictions than our CVC method. The main reason is that the two sub-nets of the CCR method are differently initialized and the two sub-nets may step into the collapse. Therefore, the sub-nets might be affected by the confirmation bias issue when using pseudo-labelling to transfer knowledge with each other. In contrast, our CVC method can efficiently avoid the situation, leading to better performance.

Parameter Analysis. We further verify the importance of our CPL method where we vary the weight of the consistency loss ω_c supervised by the pseudo-labels generated by conflicting and confident (CC) predictions. We can observe that a higher weight ω_c , *i.e.*, 2.0, can lead to a better cognitive ability of the model than the baseline, verifying our hypothesis that learning from the conflicting predictions can guarantee that the sub-nets can make consistent predictions, and thus stabilizing the training. However, if the weight is too high, the model will learn too much from self-supervision, which might be influenced by the confirmation bias issue. Thus the training might be misled and the performance might be degraded.

4.5. Qualitative results

We show the qualitative results when using different components of our method, as shown in Fig. 4. All the results are implemented on the original Pascal VOC 2012 dataset under the partition of 1/4, with ResNet-101 as the backbone of DeepLabv3+. We can see the supervised baseline (b) is prone to generate noisy predictions, the cross-consistency regularization method (c) might fail to recognize some illegible parts while our CVC method (d) and CCVC method (e) can easily recognise those hard-to-distinguish pixels, like the rear wheel of a track bike and even some small objects.

We further visualize the evolution process during training using our CCVC approach to validate the effectiveness of our proposed conflict-based pseudo-labelling (CPL) method in Fig. 5. All the results are implemented on the

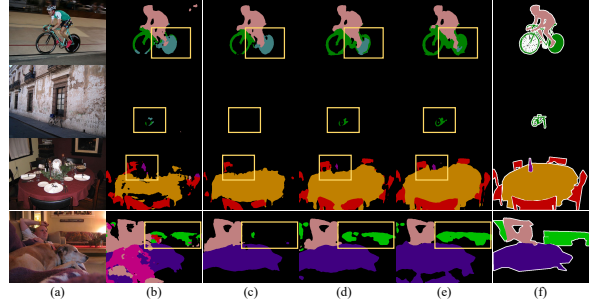


Figure 4. Qualitative results of our method from the original Pascal VOC 2012 dataset under the partition of 1/4. (a) input images, (b) the results of the supervised baseline, (c) the results of the cross-consistency regularization (CCR) method, (d) the results of our CVC method, (e) the results of our CCVC method, (f) the ground truth labels.

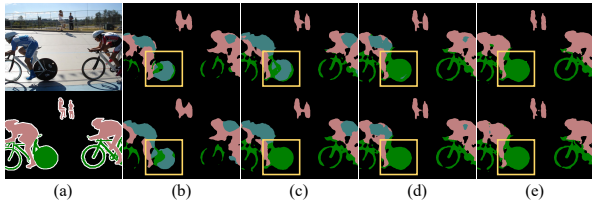


Figure 5. Qualitative results of our method on the original Pascal VOC 2012 dataset under the partition of 1/4, in which we use ResNet-101 as the backbone of our DeepLabv3+. (a) the input image and the ground-truth label, (b) the prediction results of the two sub-nets when the model is trained for 5 epochs, (c) the prediction results of the two sub-nets when the model is trained for 10 epochs, (d) the prediction results of the two sub-nets when the model is trained for 20 epochs, (e) the prediction results of the two sub-nets when the model is trained for 40 epochs. We can observe that our conflict-based pseudo-labelling (CPL) method can prevent the two sub-nets from making inconsistent predictions, which guarantees the reliability of the prediction results.

original Pascal VOC 2012 dataset under the partition of 1/4, with ResNet-101 as the backbone of DeepLabv3+. We can observe that the conflict predictions within the yellow box gradually become consistent during the training progress, indicating the effectiveness of our CPL method.

5. Conclusion

In this work, we have presented a semi-supervised semantic segmentation method based on a co-training framework, where we introduce a cross-view consistency strategy to force the two sub-nets to learn to reason the same input from irrelevant views and then exchange information with each other to generate consistent predictions. Therefore, our method can efficiently reduce the collapse and enlarge the perception of the network to produce more reliable predictions and further reduce the confirmation bias problem. Extensive experiments on the benchmark datasets have validated the effectiveness of our newly proposed approach.

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4981–4990, 2018. **1**
- [2] Inigo Alonso, Alberto Sabater, David Ferstl, Luis Montesano, and Ana C Murillo. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8219–8228, 2021. **1**
- [3] Daigang Cai, Lichen Zhao, Jing Zhang, Lu Sheng, and Dong Xu. 3djcg: A unified framework for joint dense captioning and visual grounding on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16464–16473, 2022. **1**
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. **2**
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014. **2**
- [6] 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. **2**
- [7] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. **2**
- [8] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. **2, 5**
- [9] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. **4**
- [10] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2613–2622, 2021. **2, 3, 6, 7**
- [11] Zhenghao Chen, Shuhang Gu, Guo Lu, and Dong Xu. Exploiting intra-slice and inter-slice redundancy for learning-based lossless volumetric image compression. *IEEE Transactions on Image Processing*, 31:1697–1707, 2022. **1**
- [12] Zhenghao Chen, Guo Lu, Zhihao Hu, Shan Liu, Wei Jiang, and Dong Xu. Lsvc: a learning-based stereo video compression framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6073–6082, 2022. **1**
- [13] 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. **2, 5**
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. **4, 5**
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. **2**
- [16] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. **2, 5**
- [17] Geoff French, Timo Aila, Samuli Laine, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, high-dimensional perturbations. 2019. **6**
- [18] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3146–3154, 2019. **1**
- [19] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. **4**
- [20] Dayan Guan, Jiaying Huang, Aoran Xiao, and Shijian Lu. Unbiased subclass regularization for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9968–9978, 2022. **7**
- [21] Meng-Hao Guo, Cheng-Ze Lu, Qibin Hou, Zhengning Liu, Ming-Ming Cheng, and Shi-Min Hu. Segnext: Rethinking convolutional attention design for semantic segmentation. *arXiv preprint arXiv:2209.08575*, 2022. **2**
- [22] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *2011 international conference on computer vision*, pages 991–998. IEEE, 2011. **5**
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **4, 5**
- [24] Mostafa S Ibrahim, Arash Vahdat, Mani Ranjbar, and William G Macready. Semi-supervised semantic image segmentation with self-correcting networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12715–12725, 2020. **3**

- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 5
- [26] Zhanghan Ke, Di Qiu, Kaican Li, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pages 429–445. Springer, 2020. 7
- [27] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2017. 1
- [28] Donghyeon Kwon and Suha Kwak. Semi-supervised semantic segmentation with error localization network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9957–9967, 2022. 2, 3, 6, 7
- [29] Xin Lai, Zhuotao Tian, Li Jiang, Shu Liu, Hengshuang Zhao, Liwei Wang, and Jiaya Jia. Semi-supervised semantic segmentation with directional context-aware consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1205–1214, 2021. 3, 6
- [30] Jungbeom Lee, Jooyoung Choi, Jisoo Mok, and Sungroh Yoon. Reducing information bottleneck for weakly supervised semantic segmentation. *Advances in Neural Information Processing Systems*, 34:27408–27421, 2021. 1
- [31] Xiangtai Li, Houlong Zhao, Lei Han, Yunhai Tong, Shaohua Tan, and Kuiyuan Yang. Gated fully fusion for semantic segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11418–11425, 2020. 1
- [32] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022. 1
- [33] Jiaheng Liu, Jinyang Guo, and Dong Xu. Apsnet: Toward adaptive point sampling for efficient 3d action recognition. *IEEE Transactions on Image Processing*, 31:5287–5302, 2022. 1
- [34] Shikun Liu, Shuaifeng Zhi, Edward Johns, and Andrew J Davison. Bootstrapping semantic segmentation with regional contrast. *arXiv preprint arXiv:2104.04465*, 2021. 1, 6
- [35] Yuyuan Liu, Yu Tian, Yuanhong Chen, Fengbei Liu, Vasileios Belagiannis, and Gustavo Carneiro. Perturbed and strict mean teachers for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4258–4267, 2022. 2, 3, 6, 7
- [36] 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. 2
- [37] Robert Mendel, Luis Antonio de Souza, David Rauber, Joao Paulo Papa, and Christoph Palm. Semi-supervised segmentation based on error-correcting supervision. In *European Conference on Computer Vision*, pages 141–157. Springer, 2020. 1, 2, 3
- [38] Viktor Olsson, Wilhelm Tranheden, Juliano Pinto, and Lennart Svensson. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1369–1378, 2021. 2, 3
- [39] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12674–12684, 2020. 2, 3, 7
- [40] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan Yuille. Deep co-training for semi-supervised image recognition. In *Proceedings of the european conference on computer vision (eccv)*, pages 135–152, 2018. 2, 3
- [41] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 2, 3
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [43] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020. 2
- [44] Yuchao Wang, Haochen Wang, Yujun Shen, Jingjing Fei, Wei Li, Guoqiang Jin, Liwei Wu, Rui Zhao, and Xinyi Le. Semi-supervised semantic segmentation using unreliable pseudo-labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4248–4257, 2022. 3, 5, 6, 7
- [45] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. 2
- [46] Lihe Yang, Lei Qi, Litong Feng, Wayne Zhang, and Yinghuan Shi. Revisiting weak-to-strong consistency in semi-supervised semantic segmentation. *arXiv preprint arXiv:2208.09910*, 2022. 2, 5
- [47] Luyu Yang, Yan Wang, Mingfei Gao, Abhinav Shrivastava, Kilian Q Weinberger, Wei-Lun Chao, and Ser-Nam Lim. Deep co-training with task decomposition for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8906–8916, 2021. 3
- [48] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. St++: Make self-training work better for semi-supervised se-

- semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4268–4277, 2022. 2, 5, 6, 7
- [49] Yuhui Yuan, Rao Fu, Lang Huang, Weihong Lin, Chao Zhang, Xilin Chen, and Jingdong Wang. Hrformer: High-resolution vision transformer for dense predict. *Advances in Neural Information Processing Systems*, 34:7281–7293, 2021. 2
- [50] Bowen Zhang, Zhi Tian, Quan Tang, Xiangxiang Chu, Xiaolin Wei, Chunhua Shen, and Yifan Liu. Segvit: Semantic segmentation with plain vision transformers. *arXiv preprint arXiv:2210.05844*, 2022. 2
- [51] Weichen Zhang, Wanli Ouyang, Wen Li, and Dong Xu. Collaborative and adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3801–3809, 2018. 5
- [52] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. 2
- [53] Zhen Zhao, Sifan Long, Jimin Pi, Jingdong Wang, and Luping Zhou. Instance-specific and model-adaptive supervision for semi-supervised semantic segmentation. *arXiv preprint arXiv:2211.11335*, 2022. 1
- [54] Zhen Zhao, Lihe Yang, Sifan Long, Jimin Pi, Luping Zhou, and Jingdong Wang. Augmentation matters: A simple-yet-effective approach to semi-supervised semantic segmentation. *arXiv preprint arXiv:2212.04976*, 2022. 1
- [55] Zhen Zhao, Luping Zhou, Yue Duan, Lei Wang, Lei Qi, and Yinghuan Shi. Dc-ssl: Addressing mismatched class distribution in semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9757–9765, 2022. 3
- [56] Zhen Zhao, Luping Zhou, Lei Wang, Yinghuan Shi, and Yang Gao. Lassl: label-guided self-training for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9208–9216, 2022. 3
- [57] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890, 2021. 1, 2
- [58] Yuanyi Zhong, Bodi Yuan, Hong Wu, Zhiqiang Yuan, Jian Peng, and Yu-Xiong Wang. Pixel contrastive-consistent semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7273–7282, 2021. 1, 6
- [59] Yanning Zhou, Hang Xu, Wei Zhang, Bin Gao, and Pheng-Ann Heng. C3-semiseg: Contrastive semi-supervised segmentation via cross-set learning and dynamic class-balancing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7036–7045, 2021. 3
- [60] Yuliang Zou, Zizhao Zhang, Han Zhang, Chun-Liang Li, Xiao Bian, Jia-Bin Huang, and Tomas Pfister. Pseudoseg:

Designing pseudo labels for semantic segmentation. *arXiv preprint arXiv:2010.09713*, 2020. 6