

A Simple Baseline for Semi-supervised Semantic Segmentation with Strong Data Augmentation*

Jianlong Yuan¹, Yifan Liu², Chunhua Shen², Zhibin Wang¹, Hao Li¹
¹ Alibaba Group ² The University of Adelaide, Australia

Abstract

Recently, significant progress has been made on semantic segmentation. However, the success of supervised semantic segmentation typically relies on a large amount of labeled data, which is time-consuming and costly to obtain. Inspired by the success of semi-supervised learning methods for image classification, here we propose a simple yet effective semi-supervised learning framework for semantic segmentation. We demonstrate that the devil is in the details: a set of simple designs and training techniques can collectively improve the performance of semi-supervised semantic segmentation significantly. Previous works [3, 25] fail to effectively employ strong augmentation in pseudo-label learning, as the large distribution disparity caused by strong augmentation harms the batch normalization statistics. We design a new batch normalization, namely distribution-specific batch normalization (DSBN) to address this problem and show the importance of strong augmentation for semantic segmentation. Moreover, we design a self-correction loss, which is effective in terms of noise resistance. We conduct a series of ablation studies to show the effectiveness of each component. Our method achieves state-of-the-art results in the semi-supervised settings on the Cityscapes and Pascal VOC datasets.

1. Introduction

Semantic segmentation/pixel labelling is one of the core tasks in visual understanding, which is widely used in scene parsing, human body parsing, and many downstream applications. It is a per-pixel classification problem, which classifies each pixel in the image into a predefined set of categories. In the past few years, semantic segmentation methods based on deep convolution neural networks (CNNs) have made tremendous progress [7, 11]. Note that, the common prerequisite for all these successes is the availability of a massive amount of pixel-level labeled data. Unfortunately, labeling such datasets is very expensive and time-

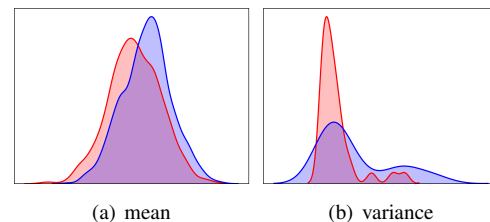


Figure 1 – Distribution mismatch of BN statistics between weak augmentation and strong augmentation. Red and blue lines denote BN statistics trained with weak and strong data augmentation, respectively.

consuming, especially in dense prediction problems, such as semantic segmentation.

Recent research reveals that semi-supervised learning (SSL), which uses a large amount of unlabeled data together with a small amount of labeled data, is greatly beneficial to classification [1, 2, 5, 6, 14, 28, 36]. These methods can be grouped into consistency methods [1, 2, 28], pseudo-labelling methods [36], representation learning [14]. A few works attempted to apply SSL to semantic segmentation. The naive student [3] uses a large model to generate the pseudo-labels with unlabeled video sequences, and applies iterative training for further improvement. They only consider the pseudo-labels with the original input images without perturbations. Recent work [25] considers adding perturbations to the images for semantic segmentation. They forward images with different perturbations in different sub-decoders and enforce the consistency between sub-decoders and the main decoder. Other work pays attention to generative methods [17, 23, 29]. AdvSemiSeg [17] and the work in [23] both employ Generative Adversarial Network (GAN) and train the model with a discriminative loss over unlabeled data and a supervised loss over labeled data.

Here, we propose an effective and efficient framework to apply SSL to semantic segmentation. We employ strong augmentation to make better use of the unlabeled data. Previous works based on consistency learning [1, 6, 35] have shown that adding noise in the procedure of learning, the pseudo-labels help improve the performance in image clas-

*JY and YL contributed equally to this work.

sification. Motivated by those works, we propose to apply strong augmentation to semantic segmentation. However, strong augmentation would inevitably affect the distribution of mean and variance in the batch normalization (BN) as shown in Figure 1. As a consequence, *one needs to be careful with computing the BN statistics in this scenario*. This may be the reason why strong augmentation was not used in semi-supervised semantic segmentation [3]. Multiple branch networks were employed to process different perturbations, as in [25]. In order to avoid the *distribution shift* caused by strong data augmentations, we propose a simple yet effective method, namely, the distribution-specific batch normalization (DSBN). We forward the strongly augmented data and the weakly-augmented data with different batch statistics during training, and inference the model with the batch statistics calculated by the weakly-augmented data.

Furthermore, as the teacher network may not perform well under all situations, some unreliable regions may be included in the generated pseudo-labels. Thus, directly learning from all the pixels will inevitably introduce label noise. Inspired by previous work in [32], which selects the unreliable image-pairs, and exchanges the prediction with the learning target to suppress the negative impact of label noise, we design a new self-correction loss (SCL) and dynamically modify the weight and learning target for each pixel for semantic segmentation.

Our approach achieves state-of-the-art performance on the Cityscapes [7] and Pascal VOC datasets [11] under a semi-supervised setting. Our main contributions are summarized as follows.

- We propose an effective and efficient semi-supervised learning framework for semantic segmentation. We employ strong augmentation during training without modifying the network structure, such as introducing sub-networks.
- A distribution-specific batch normalization is proposed to accommodate the discrepancy of data distribution in a batch caused by strong data augmentation. Besides, we design a self-training robust loss to alleviate the negative impact of label noise. The loss function is able to correct the noisy label to some extent, thus resisting noise.
- We empirically demonstrate the effectiveness of our approach, including a comparison with state-of-the-art methods, and an in-depth analysis of our approach with a detailed ablation study in the semi-supervised setting.

Next, we review some work relevant to ours.

Semi-supervised classification. Most previous works on SSL focus on classification tasks. Current state-of-the-art representative methods include consistency based methods [1, 2, 28], Pseudo labeling [36] and representation learning methods [5, 6, 14]. Consistency based methods exploit the

fact that the prediction of an unlabeled image should not change significantly with minor perturbations. UDA [35] and ReMixMatch [1] both use a weakly-augmented example to generate an artificial label and enforce consistency against strongly-augmented examples. FixMatch implements consistency training by applying strong (CutOut [10], CTAugment [1], and RandAugment [8]) and weak data augmentation to the inputs in both the unlabeled loss and the pseudo label generation, thus encouraging the outputs to be consistent for both augmented inputs. Pseudo labeling methods rely on the assumption that the pseudo-labels generated by a teacher model can benefit the training of new models. Noisy student training [36] is an iterative self-training method in this category. The correspondence among different transformations is employed for representation learning. Then a small amount of data is used to fine-tune the model.

Semi-supervised semantic segmentation. Early methods use a GAN model in semi-supervised segmentation [17, 23, 29], where an adversarial loss is trained on unlabeled data. Consistency regularization based approaches [13, 19, 21, 21, 25] have also been proposed recently, and achieve promising results. They depend on the fact that unlabeled data and labeled data have the same distribution, so it expects the trained model to have consistent and reliable predictions for both the unlabeled data and the labeled data. In these works, different unlabeled images under different transformations are used as the input to the model, and consistency loss is enforced on the prediction mask of the model. To date, self-training pushes the performance of the state-of-the-art [3, 41]. Note that none of them attempted to solve the negative impact caused by noisy pseudo-labels.

Data augmentation. Data augmentation is an effective regularization technique. The basic policy of data augmentation contains random flip, random crop, etc., which are commonly used in training vision models. In addition, the inception preprocess [30] is more sophisticated when the input image color is randomly disturbed. Recently, AutoAugment [9] improves the inception preprocess for image classification, using reinforcement learning to search for the optimal combination of augmentation strategies. RandAugment [8] proposes a significantly reduced search space which allows it to be trained on the target task and dataset, thus removing the need for a separate proxy task or dataset. All of these works tackle image classification. In this paper, we propose to apply strong augmentation to semantic segmentation.

Separated batch normalization layers. It was observed in the literature that employment of a separated batch normalization to process out-of-distribution samples can lead to performance improvement. AdvProp [34] uses an auxiliary batch normalization for accommodating domain shift coming from adversarial examples, which are proposed for

more effective adversarial training in the context of supervised image-level classification. We are inspired by this approach but employ an auxiliary batch normalization to process strongly augmented training samples which exhibit different statistics from that of the weakly augmented samples. Split-BN [38], which shares β and γ parameters, is proposed for class mismatch and image distortions of unlabeled datasets. TransNorm [31] is designed for unsupervised domain adaptation. TransNorm calculates α for a different domain. Here, our motivation is to solve the negative impact of BN statistics discrepancy caused by strong augmentation in image segmentation tasks.

Robust loss functions for training with noisy data. A few works pay attention to the loss function for learning with noisy labels, and achieve improved results. The generalized cross-entropy [39] was proposed to achieve the advantages of both mean absolute error (MAE) and categorical cross-entropy losses (CCE). Inspired by the symmetry of the Kullback-Leibler divergence, the symmetric cross-entropy [32] was proposed by combining a noise tolerance term, namely reverse cross-entropy loss, with the standard CCE loss. Nevertheless, since semantic segmentation is a pixel-level classification task, directly applying these methods to segmentation, which are designed for image-level classification, is unlikely to yield satisfactory results.

2. Our Method

2.1. Semi-supervised Semantic Segmentation

In this section, we introduce our simple semi-supervised learning framework for semantic segmentation. Different from CCT [25], which needs to design auxiliary decoders for different perturbations, our method can be applied to any existing segmentation network.

Given a small set of labeled training examples and a large set of unlabeled training examples, the small set of labeled training examples are exploited to train an initial teacher model using the standard cross-entropy loss. We then use the teacher model to generate pseudo-labels on the unlabeled images with a test-time augmentation. Following the Naive student method [3], we only generate hard pseudo-labels, *i.e.*, one-hot vectors. Then, we train a student model with *strong augmentation* to make better use of the unlabeled data. Thus, the data distribution in a batch is disturbed by the strong augmentation process, causing mismatch to that of the samples processed by weak/standard augmentation. Here, we propose the distribution-specific batch normalization (DSBN) to alleviate the negative impact of this batch distribution mismatch. Moreover, motivated by the symmetric cross entropy [32], in order to learn with potentially noisy labels, we develop a new loss function, namely, self-correction loss (SCL). SCL is used to accommodate learning with noisy labels. Finally, as in [3], we iterate the

process by putting back the student as a teacher to generate new pseudo-labels and train a new student.

Strong data augmentation for semantic segmentation (SDA). While augmentation strategies for supervised and semi-supervised image classification were extensively studied [1, 9, 28, 35], much less effort has been made for semantic segmentation. We apply random augmentation to semantic segmentation. As illustrated in Figure 2, we build a pool of operations, which contains 16 image transformation methods. At each training iteration, we randomly select a series of operations from the pool, and then combine them with standard transformations of semantic segmentation (random scale, random crop, random flip, *etc.*).

2.2. Distribution-Specific Batch Normalization

As mentioned in some previous works [4, 27, 40], batch normalization (BN) is crucial for semantic segmentation. BN computes the mini-batch mean (μ_B) and standard deviation (σ_B) [18], formulated in Equation (1). Then, as shown in Equation (2), BN uses a learnable scale (γ) and shift (β) parameter to transform the normalized distribution. Meanwhile, moving mean and moving variance are updated by Equation (3). In the testing phase, BN uses the moving mean and moving variance for whitening input activations. However, strong augmentations are likely to shift the distribution of natural images and lead to a domain gap between training images and testing images. Thus, the moving mean and moving variance calculated during the training phase would harm the testing performance. To deal with this issue, we propose the distribution-specific batch normalization (DSBN). Specifically, we have two separate BNs: BN_w for samples of weak/standard augmentation; and BN_s for

```

Transforms = ['contrast_gamma', 'contrast_linear', '
brightness', 'brightness_channel', 'equalize', '
hsv', 'invert_channel', 'blur', 'noise_gau', '
noise_pos', 'Channel_shuffle', 'dropout', '
coarse_dropout', 'multiply', 'salt_pepper', '
solarize', 'jpeg_compression']

Basic_transforms = ['random_scale', 'random_flip', '
random_crop', 'Normalization']

# data augmentation func
def SDA(N):
    """ Generate a set of distortions.
    Args:
        N: Num. of augment. transform. to apply
            sequentially
    """
    sampled_ops = np.random.choice(Transforms, N)
    return sampled_ops + Basic_transforms

```

Figure 2 – Python code for Strong Data Augmentation in Numpy.

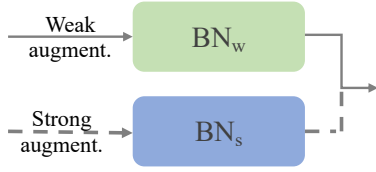


Figure 3 – Distribution-Specific Batch Normalization (DSBN). During the training phase, images with strong augmentation are passed through BN_s , while images with weak augmentation are passed through standard BN_w . During the test phase, BN_s is discarded.

samples of strong augmentation (see Figure 3). In the training phase, if a data sample comes from pseudo-labels, then BN_s is used in the forward; otherwise, BN_w is used. In the testing phase, the BN_s is deprecated, and only the BN_w is used for normalization:

$$\mu_B = \frac{1}{m} \sum_{i=1}^m x_i; \quad \sigma_B = \left[\frac{1}{m} \sum_{i=1}^m (x_i - \mu_B)^2 + \epsilon \right]^{0.5}, \quad (1)$$

$$y_i = \gamma \frac{(x_i - \mu_B)}{\sigma_B} + \beta, \quad (2)$$

$$\begin{aligned} \mu_{t+1} &= \alpha \mu_t + (1 - \alpha) \mu_{t+1}, \\ \sigma_{t+1} &= \alpha \sigma_t + (1 - \alpha) \sigma_{t+1}. \end{aligned} \quad (3)$$

2.3. Self-correction Loss

Previous work [32] has shown that cross-entropy loss masks the model overfit to noisy labels on some easier classes and under learning on some harder classes in image classification. As semantic segmentation is a per-pixel classification problem, similar challenges also exist. Moreover, in the semi-supervised learning framework, the learning targets are pseudo-labels, which may contain unreliable regions caused by the limited capacity of the current teacher model.

To avoid overfitting to the noise contained in the pseudo-labels, we propose a Self-Correction Loss (SCL) for semantic segmentation. It assigns an adaptive weight by the confidence of the network output to each pixel during the training phase. Besides, we apply a noise-robust term Reverse Cross Entropy [32] to the unreliable regions. Different from previous work, which chooses the reversed learning pairs by fixed confidence from human heuristics, we dynamically change the learning targets by comparing the output of the student network and the teacher network.

As shown in Equation (4) (i is the location index), \mathbf{y}_i^* denotes the prediction of a pixel, \mathbf{y}_i denotes the pseudo-label generated by the teacher, and w_i represents the dynamic weight which is the largest activation after softmax among all c classes (refer to Equation (5)). We use the prediction to

automatically adjust the confidence during training. If the confidence is very high, we adopt positive learning. If the confidence is very low, we believe that the pseudo-label for this pixel is unreliable. Thus, reverse learning is applied. In addition, SCL is only applied to pseudo-labels:

$$\ell_i = w_i \cdot \mathbf{y}_i \log(\mathbf{y}_i^*) + (1 - w_i) \cdot \mathbf{y}_i^* \log(\mathbf{y}_i), \quad (4)$$

$$w_i = \max\left(\frac{\exp(y_{i0}^*)}{\sum_{j=0}^c \exp(y_{ij}^*)}, \dots, \frac{\exp(y_{ic}^*)}{\sum_{j=0}^c \exp(y_{ij}^*)}\right). \quad (5)$$

3. Experiments

In this section, we first report the implementation details. Then we perform a series of ablation experiments, and analyze the results in detail. Finally, we report our results compared with other state-of-the-art methods.

3.1. Implementation Details

Training. Our implementation is built on Pytorch [26]. Following previous work [4], we use the “poly” learning rate policy, where the base learning rate is multiplied by $(1 - \text{iter}/\text{iter}_{\max})^{\text{power}}$. The initial learning rate is set to 0.01, and power is set to 0.9. We train the network using mini-batch stochastic gradient descent (SGD). We employ the crop size of 769×769 on Cityscapes with DeepLabv3Plus [4]. More details are the same as [4]. In the addition study, following [4, 20, 33, 37, 40], we adopt the mean Intersection of Union (mIoU) as the evaluation metrics.

Data augmentation. Following [4], we use mean subtraction, and apply data augmentation by random resize between 0.5 and 2 and random left-right mirror during initialization training. In the semi-supervised learning phrase, we use strong augmentation for improving the performance of self-training.

Datasets. Following [3], We conduct the main experiments and ablation studies on the Cityscapes dataset [7]. This large-scale data set contains different stereo video sequences recorded in 50 different urban street scenes. In addition to 20k weak annotation frames, there are 5,000 high-quality pixel level annotation frames, in which, 2,975; 500 and 1,525 were used for training, validation and testing. In addition, there are 20k images with rough annotations, namely train-coarse. We randomly subsample $1/8$, $1/4$, and $1/2$ of images in the standard training set to construct the pixel-level labeled data. The remaining images in the training set are used as unlabeled data.

At last, we report our results on the Pascal VOC 2012 dataset [11], which contains 21 classes including background. The standard Pascal VOC 2012 dataset has 1,449 images as the training set and 1,456 images as the validation set. We construct 1,449 images as the pixel-level labeled data. The images in the augmented set [15] (around 9k images), are used as unlabeled data.

Model	s.t.	SDA	DSBN	SCL	iter.	mIoU-1/8 (%)	mIoU-1/4 (%)	mIoU-1/2 (%)	mIoU-Full (%)
Deeplabv3plus						68.9	73.4	76.9	78.7
Deeplabv3plus	✓					71.2	76.2	78.0	79.5
Deeplabv3plus	✓	✓				71.3	76.3	78.2	79.7
Deeplabv3plus	✓	✓	✓			72.3	76.9	78.7	80.0
Deeplabv3plus	✓	✓	✓	✓		72.8	77.4	78.7	80.4
Deeplabv3plus	✓	✓	✓	✓	✓	74.1	77.8	78.7	80.5

Table 1 – Ablation study on the proposed semi-supervised learning framework. The model here is Deeplabv3plus with Xception65 backbone. ‘s.t.’ denotes self-training with pseudo-labels without strong augmentation. ‘SDA’ means the strong data augmentation. ‘DSBN’ means the distribution specify batch normalization. ‘SCL’ is the proposed self correction loss. ‘iter.’ represents iterative training. The results are evaluated on the validation set with the single-scale input. ‘mIoU-1/n’ means that we use 1/n labeled data and the remaining images in the training set are used as unlabeled data. ‘mIoU-Full’ means that we use all the labeled data in the training set, and the train-coarse are employed as the unlabeled dataset.

3.2. Ablation Study

In this subsection, we conduct the experiments to explore the effectiveness of each proposed module under different semi-supervised settings. First, we establish the baseline for our experiments. We evaluate DeepLabv3Plus based on Xception65 with cross entropy on the validation set. Following [3,36], self training adopts the baseline model as the teacher model and the teacher model generates pseudo labels on the unlabeled dataset, then the student model which is not smaller than the teacher model is trained on pseudo labels and original labeled images. We use the scales including 0.5, 0.75, 1.0, 1.5, 1.75 and flip for the remaining images to generate pseudo labels. For fair comparison and quick training, the student model is the same model initialized by the teacher model. The main results are shown in Table 1, and we will give detail discussions on each module in the following sections. All the ablation study are conducted on Cityscapes.

3.2.1 Different semi-supervised settings

We follow previous works [13, 17, 29] to divide the train set into labeled data and unlabeled data according to different proportions. We use 1/8, 1/4, and 1/2 of the training dataset as labeled data and other images as the unlabeled data. The evaluation results are shown in Table 1 as mIoU-1/8, mIoU-1/4, mIoU-1/2, respectively. From the Table 1, we can see that our algorithm can effectively improve the performance by 5.2%, 4.4%, 1.8%, 1.8%, respectively. When the total number of labeled data and unlabeled data does not change, the semi-supervised learning framework will have a larger improvement, if we have less labeled images. In the meantime, the final results will be better if we have more high-quality labeled images. We can also see that only half the amount of labeled data with the help of proposed semi-supervised learning framework can achieve the same results (mIoU of 78.7%) as we use all labeled training data with supervised learning (row 1 and column 10).

When the total number of available data increases, semi-

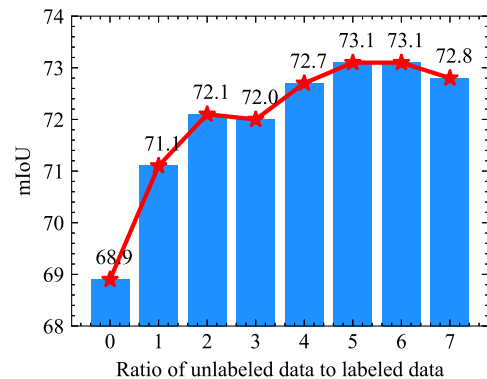


Figure 4 – Different ratios of unlabeled data to labeled data.

supervised learning will play a more valuable role. We use all high-quality labeled data in the training set as the labeled data and employ the train-coarse as unlabeled data to conduct the experiment, and the evaluation results are referred to as mIoU-Full in Table 1. We can see that under a strong baseline with large amount of labeled data, if extra unlabeled data can be obtained, the performance will be further improved with the proposed semi-supervised learning framework.

In practical applications, labeled data is often limited, but unlabeled data is easy to obtain. To explore the impact of increasing the ratio of unlabeled data in the training set. We fix the labeled data as 1/8 of the training set, and increase the number of the unlabeled data. The results are shown in Figure 4. From Figure 4, we can see that increasing ratios of the unlabeled data to the labeled data would improve the results, but the growth trend of the performance is gradually flattening limited by the capacity of the initial teacher model trained on the fixed labeled data. The performance even slightly drops if too many unlabeled images are introduced during training.

3.2.2 Impact of the strong augmentation

In this section, we will show the impact of the proposed strong augmentation and the DSBN. First, as shown in Table 1, we compare the results of the naive self-training and the self-training combined with strong augmentation and the proposed DSBN. Although self-training can improve the performance over the baseline, self training directly uses pseudo labels that makes the network difficult to learn and introduces noise, as discussed in Section 2. Meanwhile, in the case of only using strong augmentation, there is no significant improvement in performance. On the contrary, our method performs favorably against naive self training with strong augmentation and DSBN. In particular, compared with naive self training, our method improves the mIoU by 1.1%, 0.7%, 0.7% and 0.5% based on DeepLabv3Plus. In the following sections, we analyze the distribution of the BN statistics, and then we compare the strong augmentation with different BN settings. Finally, we provided some discussions on how does the strong augmentation work with ground truth labels and the pseudo labels.

Visualization of the BN statistics. To verify our statement that the distribution of the BN will be affect by the strong augmentation, we show the statistics of BN training with strong augmentation input images and the weakly augmentation input images. The visualization results can be found in Figure 1.

The distribution of the mean and variance obtained from the whole training statistics deviates from the test set after adding strong augmentation. As shown in Figure 1, we calculate the distribution of the mean and the variance of the BN with the weakly augmentation and the strong augmentation. The blue line denotes the distribution of the BN with weakly augmentation, while the red line represents the BN with strong augmentation. We can see that strong augmentation has changed the distribution of BN statistics, which may lead to a domain gap. We think this is the reason why strong augmentation has not brought much performance improvement, as shown in Table 1.

Impact of the DSBN. To show the effectiveness of proposed DSBN, we conduct experiments to compare three different BN settings as follows:

- Trainable BN: the mean and the variance are updated during training with both weak augmentation data and strong augmentation data.
- Fixed BN: the mean and the variance are fixed during training, and initialized with the pre-trained weights from the ImageNet classification.
- DSBN: different mean and the variance are updated during training for weak augmentation data and strong augmentation data, separately.

The baseline model is DeepLabV3Plus trained with 1/8 labeled data on the training set. The other 7/8 images are

used as the unlabeled dataset in the semi-supervised learning. The results are shown in Table 2. We can see that with the proposed DSBN, the performance can be improved by 0.7%, which indicates the DSBN can help to solve the negative effect caused by the strong augmentation.

Discussions. To further understand the impact of the strong augmentation and DSBN on different data groups, we apply strong augmentation to only ground truth labels and pseudo labels.

- $1/8$ ground truth. We train the model with only 1/8 images of the training dataset with corresponding ground truth labels.
- Full-ground truth. We train the model with all images of the training dataset with corresponding ground truth labels.
- $7/8$ pseudo labels. We train the model with 7/8 images of the training dataset with corresponding pseudo labels.

For each settings, we apply weekly augmentation, strong augmentation and DSBN during training. The experiment results are shown in Table 3. We can see that under the fully supervised settings, the strong augmentation can not bring extra improvements. Meanwhile, when applying to the unlabeled data with pseudo labels, the mIOU improved by 0.8%. These observations indicate that improvements under semi-supervised settings are not come from extra image transformation types on labeled data, but benefit from make further use of unlabeled data with strong augmentations. From Table 3, we also find that the DSBN can contribute to the training. *In addition, to our knowledge, this is the first time to apply the strong augmentation to semantic segmentation under fully-supervised setting, and improvement is observed.*

3.2.3 Impact of the self-correction loss

To further improve the performance and reduce the impact of noise, we employ a self correction loss (SCL). From Table 1, we can see that the proposed SCL will play a more important role when the pseudo labels contain more noise. The teacher model trained with a limited labeled dataset may get less useful knowledge, and can not generalize well on unlabeled dataset. Thus, the level of the noise is higher than that

Baseline	mIoU (%)
ST	71.2
ST + SDA + Trainable BN	71.3
ST + SDA + Fixed BN	71.6
ST + SDA + DSBN	72.3

Table 2 – Ablation study for different batch normalization methods.

Methods	mIoU (%)
Training with ground-truth labels	
1/8-ground truth	68.9
1/8-ground truth + SDA	68.6
1/8-ground truth + SDA + DSBN	69.5
Full-ground truth	78.7
Full-ground truth + SDA	78.7
Full-ground truth + SDA + DSBN	79.2
Training with pseudo-labels	
7/8-pseudo labels	70.6
7/8-pseudo labels + SDA	71.4
7/8-pseudo labels + SDA + DSBN	72.5

Table 3 – Effect of the strong augmentation on different data. We apply the strong augmentation to the baselines trained with only labeled data, and only unlabeled data with pseudo-labels, separately.

training on more labeled data. When we train with 1/8 or 1/4 labeled images, the SCL can improve the mIoU by 0.5%.

Furthermore, we show the mIoU during the training phase with and without the proposed SCL in Figure 5. When training without the SCL, the result will be divergent quickly. On the contrary, the result is relatively stable.

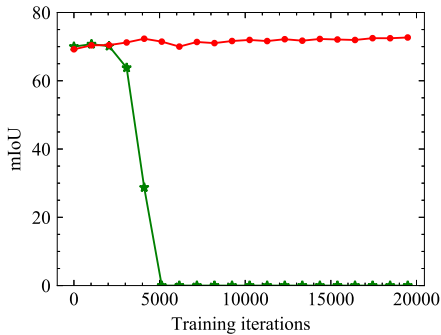


Figure 5 – This curve shows the validation mIoU w.r.t. the number of iterations during training. The green line represents the SCE loss [32]. In addition, we set the same ratio between CE and RCE. The red line represents our proposed loss with the same initial weights as the SCE loss.

3.2.4 Impact of iterative training

Following Naive Students [3], we also test the impact of the iterative training in our semi-supervised learning framework. Table 1 has shown that the iteration training will become less useful if more labeled data can be obtained. Because the initial teacher model trained with more labeled data can already produce satisfactory pseudo labels. We show the change of the performance for the number of iterations in Figure 6. The baseline model is DeepLabV3Plus

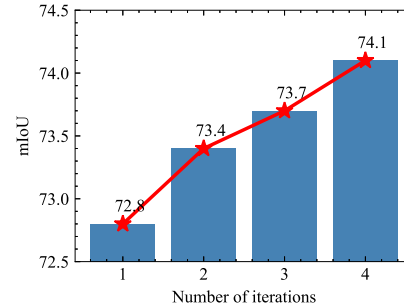


Figure 6 – The performance of the DeepLabV3Plus under different number of iterations.

Method	Model	1/8	1/4	1/2	Full
AdvSemiSeg [17]	DeepLabV2	58.8	62.3	65.7	66.0
S4GAN + MLMT [23]	DeepLabV2	59.3	61.9	-	65.8
CutMix [13]	DeepLabV2	60.3	63.87	-	67.7
DST-CBC [12]	DeepLabV2	60.5	64.4	-	66.9
ClassMix [24]	DeepLabV2	61.4	63.6	66.3	66.2
ECS [22]	DeepLabV3Plus	67.4	70.7	72.9	74.8
Baseline	DeepLabV2	60.6	66.7	69.3	70.1
Ours	DeepLabV2	67.6	69.3	70.7	70.1
Baseline	DeepLabV3Plus	68.9	73.4	76.9	78.7
Ours	DeepLabV3Plus	74.1	77.8	78.7	78.7

Table 4 – Compared with state-of-the-art methods on the Cityscapes val set. Here ‘1/n’ means that we use 1/n labeled data and the remaining images in the training set are used as unlabeled data.

trained with 1/8 labeled images. We can see that more iterations indeed help improve the performance, but the growth trend of the performance is gradually flattening.

Visualization results Several visualization results are shown in Figure 7. We can see that the baseline and self-training can not well separate the objects (especially large-size objects such as bus, truck, train, sidewalk) completely while ours corrects these errors. Our method also performs well on these small-size object classes, such as pole and bicycle, compared to the baseline model and self-training.

3.3. Comparison with State-of-the-art Methods

Cityscapes. We conduct the comparison experiments with other state-of-the-art algorithms on the Table 4. To make a fair comparison, we also apply our proposed framework to DeepLabV2 following [17]. In particular, DeepLabV2 is based on ResNet-101 initialized with the pre-trained weights on ImageNet classification. The proposed framework achieves 67.6%, 69.3%, 70.7% with DeepLabV2 and 74.1%, 77.8%, 78.7% with DeepLabV3Plus, respectively, which performs favorably against previous state-of-the-art methods. More details of ablation studies of with DeepLabV2 can be found in the supplementary materials.

Pascal VOC. We follow DeepLabV3Plus [4] training details with the 513×513 crop size on Pascal VOC [11]. Ta-

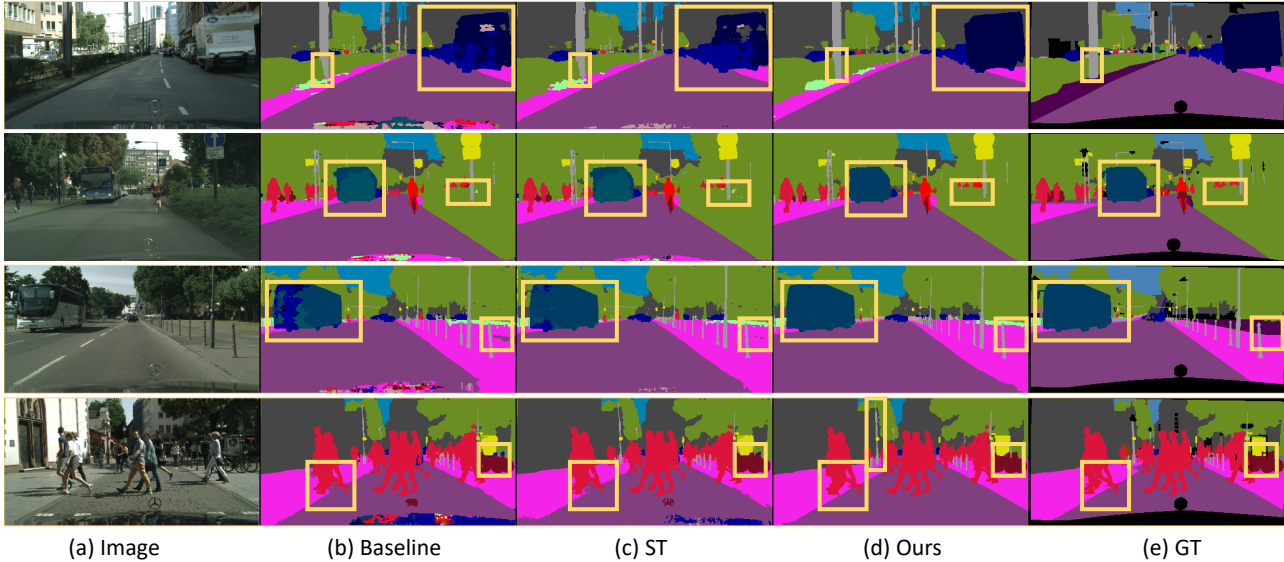


Figure 7 – Qualitative results on the Cityscapes dataset. The baseline method is trained with $1/8$ of all labeled images. ‘ST’ represents the self-training trained with $7/8$ pseudo-labels without our methods. ‘Ours’ is our framework with $1/8$ labelled data. The proposed semi-supervised approach produces improved results compared to the baseline and the naive self training. We highlight the details in yellow boxes.

Method	Model	mIoU
GANSeg [29]	VGG16	64.1
AdvSemSeg [17]	ResNet-101	68.4
CCT [25]	ResNet-50	69.4
PseudoSeg [42]	ResNet-101	73.2
Ours	ResNet-101	75.0
Ours	Xception-65	79.3
Fully supervised	ResNet-101	78.3
Fully supervised	Xception-65	79.2

Table 5 – Comparison with state-of-the-art on the Pascal VOC val set (w/ unlabeled data). We use the official training set (1.4k) as labeled data, and the augmented set (9k) as unlabeled data.

Table 5 lists the performance results of other state-of-the-art methods and ours. We adopt the single-scale testing on our experiments. In order to make a fair comparison, we did experiments under different backbones. Our framework achieves 75.0% mIoU based on Resnet-101 [16] and 79.3% mIoU based on Xception-65 [4], which outperforms the PseudoSeg [42] by 1.8%. At the same time, we find that our results based on Xception-65 can achieve the same performance as that of the fully supervised setting. In addition, we find that the improvement of iterative training on Pascal VOC is limited, which might be due to the small number of samples.

4. Conclusion

In this work, we construct a simple semi-supervised learning framework for semantic segmentation. It employs

strong augmentation with distribution-specific batch normalization, together with pseudo-label learning. In particular, DSBN can effectively avoid the BN statistics shift caused by strong augmentation. Meanwhile, we have also designed a self-correction loss, which effectively alleviates the label noises introduced by pseudo labeling. Quantitative and qualitative comparisons show that the proposed method performs favorably against recent state-of-the-art semi-supervised approaches.

References

- [1] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remix-match: Semi-supervised learning with distribution matching and augmentation anchoring. In *Proc. Int. Conf. Learn. Represent.*, 2019. 1, 2, 3
- [2] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 5049–5059, 2019. 1, 2
- [3] Liang-Chieh Chen, Raphael Gontijo Lopes, Bowen Cheng, Maxwell D Collins, Ekin D Cubuk, Barret Zoph, Hartwig Adam, and Jonathon Shlens. Semi-supervised learning in video sequences for urban scene segmentation. *Proc. Eur. Conf. Comput. Vis.*, 2020. 1, 2, 3, 4, 5, 7
- [4] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proc. Eur. Conf. Comput. Vis.*, pages 801–818, 2018. 3, 4, 7, 8
- [5] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised mod-

- els are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020. 1, 2
- [6] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 1, 2
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3213–3223, 2016. 1, 2, 4
- [8] Ekin Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proc. Adv. Neural Inform. Process. Syst.*, 2020. 2
- [9] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V. Le. AutoAugment: Learning augmentation strategies from data. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 113–123, 2019. 2, 3
- [10] Terrance DeVries and Graham W. Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 2
- [11] Mark Everingham, Luc Van Gool, Christopher Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010. 1, 2, 4, 7
- [12] Zhengyang Feng, Qianyu Zhou, Guangliang Cheng, Xin Tan, Jianping Shi, and Lizhuang Ma. Semi-supervised semantic segmentation via dynamic self-training and class-balanced curriculum. *arXiv preprint arXiv:2004.08514*, 2020. 7
- [13] Geoff French, Timo Aila, Samuli Laine, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, high-dimensional perturbations. *Proc. Brit. Mach. Vis. Conf.*, 2019. 2, 5, 7
- [14] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. In *Proc. Adv. Neural Inform. Process. Syst.*, 2020. 1, 2
- [15] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proc. Int. Conf. Comput. Vis.*, pages 991–998. IEEE, 2011. 4
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2016. 8
- [17] Wei-Chih Hung, Yi-Hsuan Tsai, Yan-Ting Liou, Yen-Yu Lin, and Ming-Hsuan Yang. Adversarial learning for semi-supervised semantic segmentation. *Proc. Brit. Mach. Vis. Conf.*, 2018. 1, 2, 5, 7, 8
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Proc. Int. Conf. Mach. Learn.*, 2015. 3
- [19] Zhanghan Ke, Di Qiu, Kaican Li, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. *arXiv preprint arXiv:2008.05258*, 2020. 2
- [20] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1925–1934, 2017. 4
- [21] Wenfeng Luo and Meng Yang. Semi-supervised semantic segmentation via strong-weak dual-branch network. In *Proc. Eur. Conf. Comput. Vis.*, pages 784–800. Springer, 2020. 2
- [22] Robert Mendel, Luis Antonio de Souza, David Rauber, João Paulo Papa, and Christoph Palm. Semi-supervised segmentation based on error-correcting supervision. In *Proc. Eur. Conf. Comput. Vis.*, pages 141–157. Springer, 2020. 7
- [23] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high-and low-level consistency. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019. 1, 2, 7
- [24] Viktor Olsson, Wilhelm Tranheden, Juliano Pinto, and Lennart Svensson. ClassMix: Segmentation-based data augmentation for semi-supervised learning. In *Proc. Winter Conf. Appl. Comput. Vis.*, 2021. 7
- [25] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12674–12684, 2020. 1, 2, 3, 8
- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 8026–8037, 2019. 4
- [27] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. In-place activated batchnorm for memory-optimized training of dnns. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5639–5647, 2018. 3
- [28] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. FixMatch: Simplifying semi-supervised learning with consistency and confidence. In *Proc. Adv. Neural Inform. Process. Syst.*, 2020. 1, 2, 3
- [29] Nasim Souly, Concetto Spampinato, and Mubarak Shah. Semi supervised semantic segmentation using generative adversarial network. In *Proc. Int. Conf. Comput. Vis.*, pages 5688–5696, 2017. 1, 2, 5, 8
- [30] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1–9, 2015. 2
- [31] Ximei Wang, Ying Jin, Mingsheng Long, Jianmin Wang, and Michael Jordan. Transferable normalization: Towards improving transferability of deep neural networks. *Proc. Adv. Neural Inform. Process. Syst.*, 2019. 3
- [32] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 322–330, 2019. 2, 3, 4, 7

- [33] Zifeng Wu, Chunhua Shen, and Anton van den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90:119–133, 2019. 4
- [34] Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan L. Yuille, and Quoc V. Le. Adversarial examples improve image recognition. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 819–828, 2020. 2
- [35] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V. Le. Unsupervised data augmentation for consistency training. In *Proc. Adv. Neural Inform. Process. Syst.*, 2020. 1, 2, 3
- [36] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V. Le. Self-training with noisy student improves imagenet classification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10687–10698, 2020. 1, 2, 5
- [37] Jianlong Yuan, Zelu Deng, Shu Wang, and Zhenbo Luo. Multi receptive field network for semantic segmentation. In *Proc. Winter Conf. Appl. Comput. Vis.*, pages 1883–1892. IEEE, 2020. 4
- [38] Michal Zajkac, Konrad Zolna, and Stanislaw Jastrzebski. Split batch normalization: Improving semi-supervised learning under domain shift. In *Proc. Int. Conf. Learn. Represent. 2nd Learning from Limited Labeled Data Workshop*, 2019. 3
- [39] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 8778–8788, 2018. 3
- [40] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2881–2890, 2017. 3, 4
- [41] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin Cubuk, and Quoc V. Le. Rethinking pre-training and self-training. In *Proc. Adv. Neural Inform. Process. Syst.*, 2020. 2
- [42] Yuliang Zou, Zizhao Zhang, Han Zhang, Chun-Liang Li, Xiao Bian, Jia-Bin Huang, and Tomas Pfister. PseudoSeg: Designing pseudo labels for semantic segmentation. In *Proc. Int. Conf. Learn. Represent.*, 2021. 8