

# Augmentation Matters: A Simple-yet-Effective Approach to Semi-supervised Semantic Segmentation

Zhen Zhao<sup>1</sup> Lihe Yang<sup>3</sup> Sifan Long<sup>4</sup> Jimin Pi<sup>2</sup> Luping Zhou<sup>1\*</sup> Jingdong Wang<sup>2\*</sup>  
<sup>1</sup>University of Sydney      <sup>2</sup>Baidu Inc.      <sup>3</sup>Nanjing University      <sup>4</sup>Jilin University

## Abstract

Recent studies on semi-supervised semantic segmentation (SSS) have seen fast progress. Despite their promising performance, current state-of-the-art methods tend to increasingly complex designs at the cost of introducing more network components and additional training procedures. Differently, in this work, we follow a standard teacher-student framework and propose **AugSeg**, a simple and clean approach that focuses mainly on data perturbations to boost the SSS performance. We argue that various data augmentations should be adjusted to better adapt to the semi-supervised scenarios instead of directly applying these techniques from supervised learning. Specifically, we adopt a simplified intensity-based augmentation that selects a random number of data transformations with uniformly sampling distortion strengths from a continuous space. Based on the estimated confidence of the model on different unlabeled samples, we also randomly inject labelled information to augment the unlabeled samples in an adaptive manner. Without bells and whistles, our simple AugSeg can readily achieve new state-of-the-art performance on SSS benchmarks under different partition protocols<sup>1</sup>.

## 1. Introduction

Supervised semantic segmentation studies [5, 6, 37, 53] have recently achieved tremendous progress, but their success depends closely on large datasets with high-quality pixel-level annotations. Delicate and dense pixel-level labelling is costly and time-consuming, which becomes a significant bottleneck in practical applications with limited labelled data. To this end, semi-supervised semantic segmentation (SSS) [27, 39] has been proposed to train models on less labelled but larger amounts of unlabeled data.

Consistency regularization [42, 43], the currently dominant fundamental SSS method, effectively incorporates the

\*Corresponding authors (luping.zhou@sydney.edu.au, wangjingdong@baidu.com). This work is supported by Australian Research Council (ARC DP200103223).

<sup>1</sup>Code and logs: <https://github.com/zhenzhao/AugSeg>.




Figure 1. Comparison between current SOTAs and our simple AugSeg on Pascal VOC 2012, using R101 as the encoder.

training on unlabeled data into standard supervised learning [16, 44]. It relies on the label-preserving data or model perturbations to produce the prediction disagreement on the same inputs, such that unlabeled samples can be leveraged to train models even if their labeled information is unknown. Some studies in [17, 29, 50, 51] explored different data augmentations to benefit the SSS training while works in [7, 16, 46] mainly focused on various model perturbations to obtain competitive SSS performance. On top of these fundamental designs, recent state-of-the-art (SOTA) methods aim to integrate extra auxiliary tasks [1, 47, 56, 57], *e.g.*, advanced contrastive learning techniques, and more trainable modules [28, 30, 36, 38], *e.g.* multiple ensemble models and additional correcting networks, to further improve the SSS performance. Despite their promising performance, SSS studies along this line come at the cost of requiring more complex methods, *e.g.*, extra network components or additional training procedures.

In this paper, we break the trend of recent SOTAs that combine increasingly complex techniques and propose **AugSeg**, a simple-yet-effective method that focuses mainly on data perturbations to boost the SSS performance. Although various auto data augmentations [9, 10] and cutmix-related transformations [17, 52] in supervised learning have

Method	Augmentations		More Supervision			Pseudo-rectifying		
	SDA	FT	MBSL	CT	UCL	UAFS	ACN	PR
CCT [44]		✓	✓	✓				
ECS [38]			✓				✓	
SSMT [26]	✓		✓			✓		
PseudoSeg [58]	✓					✓		
CAC [31]			✓		✓	✓		
DARS [24]	✓		✓					✓
PC <sup>2</sup> Seg [56]	✓		✓		✓	✓		
C3-Semiseg [57]	✓				✓	✓		✓
ReCo [34]	✓			✓	✓	✓		
CPS [7]	✓			✓				
ST++ [50]	✓		✓					
ELN [30]	✓		✓				✓	
USRN [20]	✓		✓			✓		✓
PSMT [36]	✓	✓	✓			✓		
U <sup>2</sup> PL [47]	✓				✓	✓		
<b>AugSeg (ours)</b>	✓							

Table 1. Comparison of recent SSS algorithms in terms of “Augmentations”, “More supervision”, and “Pseudo-rectifying” (sorted by their publication date). We explain the abbreviations as follows. “SDA”: Strong data augmentations, including various intensity-based and cutmix-related augmentations, “FT”: Feature-based augmentations, “MBSL”: multiple branches, training stages, or losses, “CT”: Co-training, “UCL”: unsupervised contrastive learning, “UAFS”: uncertainty/attention filtering/sampling, “ACN”: additional correcting networks, “PR”: prior-based re-balancing techniques. **Note that**, branches of “more supervision” and “pseudo-rectifying” typically require more training efforts. Differently, our method enjoys the best simplicity but the highest performance.

been extensively utilized in previous SSS studies, we argue that these augmentations should be adjusted precisely to better adapt to the semi-supervised training. **On one hand**, these widely-adopted auto augmentations are essentially designed for supervised paradigm and aim to search the optimal augmentation strategies from a predefined finite discrete space. Their optimal objective is constant and clear across the training course. However, data perturbations in semi-supervised learning consist in generating prediction disagreement on the same inputs, without a constant and specific objective or a predefined discrete searching space. Thus, we simplify existing randomAug [10] and design a highly random intensity-based augmentation, which selects a random number of different intensity-based augmentations and a random distortion strength from a continuous space. **On the other hand**, random copy-paste [18] among different unlabeled samples can yield effective data perturbations in SSS, but their mixing between corresponding pseudo-labels can inevitably introduce confirmation bias [3], especially on these instances with less confident predictions of the model. Considering the utilization efficiency of unlabeled data, we simply mix labeled samples with these less confident unlabeled samples in a random and adaptive manner, *i.e.*, adaptively injecting labeled information to stabilize the training on unlabeled data. Benefiting from the simply random and collaborative designs, AugSeg

requires no extra operations to handle the distribution issues, as discussed in [51].

Despite its simplicity, AugSeg obtains new SOTA performance on popular SSS benchmarks under various partition protocols. As shown in Figure 1, AugSeg can consistently outperform current SOTA methods by a large margin. For example, AugSeg achieves a high mean intersection-over-union (mIoU) of 75.45% on classic Pascal VOC 2012 using only 183 labels compared to the supervised baseline of 59.10% and previous SOTA of 71.0% in [50]. We attribute these remarkable performance gains to our revision – that various data augmentations are simplified and adjusted to better adapt to the semi-supervised scenarios. Our main contributions are summarized as follows,

- We break the trend of SSS studies that integrate increasingly complex designs and propose AugSeg, a standard and simple two-branch teacher-student method that can achieve readily better performance.
- We simply revise the widely-adopted data augmentations to better adapt to SSS tasks by injecting labeled information adaptively and simplifying the standard RandomAug with a highly random design.
- We provide a simple yet strong baseline for future SSS studies. Extensive experiments and ablations studies are conducted to demonstrate its effectiveness.

## 2. Related work

The key to semi-supervised learning lies in effectively leveraging the unlabeled data [19,21,33,42]. Recent consistency regularization (CR) [32,46] has become a fundamental semi-supervised technique to train models on labeled and unlabeled data simultaneously. Such CR-based methods, in either classification [14,45,48,54] or segmentation tasks [17,25,36,58], rely on various perturbation techniques to generate disagreement on the same inputs, so that models can be trained by enforcing prediction consistency on unlabeled data without knowing labeled information. Along this line, many SSS methods have been proposed recently.

Based on our summary, as shown in Table 1, there are three different main directions to enhance the SSS performance, including “augmentations”, “more supervision,” and “pseudo-rectifying”. Almost all existing studies applied various strong data augmentations to perturb unlabeled data while some of them [36,44,49] also perturbed the inputs at the feature level. In the branch of “more supervision”, multiple training branches, training stages, or losses (MBSL) are widely adopted from the perspective of model perturbations [26,31,36,50]. As the quality of pseudo-labels is critical for semi-supervised training [55], ECS [38] and ELN [30] also introduced additional trainable correcting

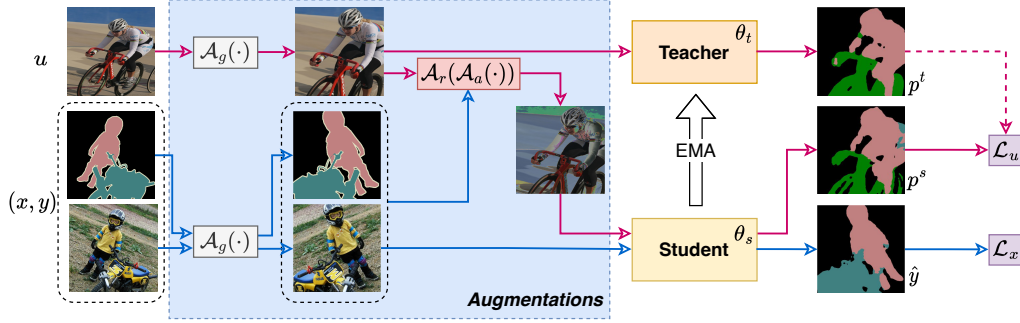


Figure 2. Diagram of AugSeg. In a standard teacher-student framework, AugSeg trains the student model, parameterized by  $\theta_s$ , on labeled data  $(x, y)$  and unlabeled data  $u$  simultaneously, via minimizing the corresponding supervised loss  $\mathcal{L}_x$  and unsupervised consistency loss  $\mathcal{L}_u$ , respectively. The teacher model, parameterized by  $\theta_t$ , is updated by the exponential moving averaging (EMA) of  $\theta_s$ , and generates the pseudo-label on unlabeled data,  $p^t$ . The core of AugSeg is to apply various augmentation techniques on input unlabeled samples, including the weak geometrical augmentation  $\mathcal{A}_g$ , the random intensity-based augmentation  $\mathcal{A}_r$  and the adaptive label-aided augmentation  $\mathcal{A}_a$ . The red and blue lines represent the forward path of labeled and unlabeled data, respectively. The dashed line means “stop gradient”.

networks (ACN) to further polish the pseudo-labels. Recent SOTA methods [20, 30, 36, 47] can achieve promising performance at the cost of combining increasingly complex mechanisms, *e.g.* contrastive learning [35] and multiple ensembling models. Differently, in this work, we aim to propose a simple and clean method that relies mainly on data augmentations to boost the SSS performance.

As the most straightforward and effective way to produce label-preserving perturbations, data augmentations have played a central role in CR-based semi-supervised studies [17, 40, 48]. Recently, various auto augmentation strategies [9, 11, 41] in supervised learning have been widely adopted in semi-supervised research. However, directly applying such techniques is not satisfying for the following reasons. The goal of auto augmentations is to find out the optimal augmentation strategies. Such studies aim to search best augmentation operations and distortion strengths from a pre-defined finite discrete space. In contrast, the objective of applying data augmentation in SSS is to generate different inputs without specific goals and searching spaces. Besides, as discussed in [51], directly applying these augmentations may over-distort unlabeled data and hurt the data distribution, resulting in performance degradation. Instead of using additional rectifying strategies like distribution-specific BN [4], we simplify the standard randomAug [10] with a highly random design. Instead of using a predefined number of augmentations with finite and discrete strength possibilities, we select a random number of augmentations and sample the augmentation strength uniformly from a continuous interval. In this way, our design enjoys better data diversity and is less likely to over-distort samples.

### 3. Augseg

In this section, we first present an overview of our simple AugSeg and then describe our main augmentation strate-

gies, random intensity-based and adaptive cutmix-based augmentations, in Section 3.2 and Section 3.3, respectively.

#### 3.1. Overview

Following current dominant consistency regularization methods in semi-supervised learning [43], we train our segmentation model on labeled and unlabeled data simultaneously. As shown in Figure 2, we adopt a simple and clean semi-supervised framework, which consists of a student model and a teacher model, parameterized by  $\theta_s$  and  $\theta_t$ , respectively. Specifically, the teacher model is capable of producing pseudo-labels for training on unlabeled data, and will be updated gradually via the exponential moving averaging of the student weights, *i.e.*,

$$\theta_t \leftarrow \alpha \theta_t + (1 - \alpha) \theta_s, \quad (1)$$

where  $\alpha$  is a common momentum parameter, which is, following [46], set as 0.999 by default. On the other hand, at each iteration, provided with a batch of labeled samples  $\mathcal{B}_x = \{(x_i, y_i)\}_{i=1}^{|\mathcal{B}_x|}$  and a batch of unlabeled samples  $\mathcal{B}_u = \{u_i\}_{i=1}^{|\mathcal{B}_u|}$ , we aim to train the student model via minimizing a supervised loss  $\mathcal{L}_x$  and an unsupervised consistency loss  $\mathcal{L}_u$  at the same time. Thus the total training loss for the student model is,

$$\mathcal{L} = \mathcal{L}_x + \lambda_u \mathcal{L}_u, \quad (2)$$

where  $\lambda_u$  is a scalar hyper-parameter to adjust the unsupervised loss weight. Similar to most SSS methods [27, 38, 44], we adopt a standard pixel-wise cross-entropy loss  $\ell_{ce}$  to train on labeled data directly,

$$\mathcal{L}_x = \frac{1}{|\mathcal{B}_x|} \sum_{i=1}^{|\mathcal{B}_x|} \frac{1}{H \times W} \sum_{j=1}^{H \times W} \ell_{ce}(\hat{y}_i(j), y_i(j)), \quad (3)$$

Weak Geometrical Augmentation - Apply all	
Random Scale	Randomly resizes the image by [0.5, 2.0].
Random Flip	Horizontally flips the image with a probability of 0.5.
Random Crop	Randomly crops an region from the image.
Random Intensity-based Augmentation - Apply $k$ randomly	
Identity	Returns the original image.
Autocontrast	Maximizes (normalize) the image contrast.
Equalize	Equalize the image histogram.
Gaussian blur	Blurs the image with a Gaussian kernel.
Contrast	Adjusts the contrast of the image by [0.05, 0.95].
Sharpness	Adjusts the sharpness of the image by [0.05, 0.95].
Color	Enhances the color balance of the image by [0.05, 0.95]
Brightness	Adjusts the brightness of the image by [0.05, 0.95]
Hue	Jitters the hue of the image by [0.0, 0.5]
Posterize	Reduces each pixel to [4,8] bits.
Solarize	Inverts image pixels above a threshold from [1,256).

Table 2. List of various image transformations in the weak geometrical augmentation and random intensity-based augmentation.

where  $\hat{y}_i = f(\mathcal{A}_g(x_i); \theta_s)$ , represents the segmentation result of the student model on the  $i$ -th weakly-augmented labeled instance.  $j$  represents the  $j$ -th pixel on the image or the corresponding segmentation mask with a resolution of  $H \times W$ . The weak geometrical augmentation  $\mathcal{A}_g$ , as shown in Table 2, includes standard resizing, cropping, and flipping operations. As for leveraging the unlabeled data, which is the key to semi-supervised learning, we rely mainly on the data perturbation  $\mathcal{T}(\cdot)$  to generate the prediction disagreement. First, we obtain the segmentation predictions,  $p_i^s$  and  $p_i^t$ , of the student model on augmented  $\mathcal{T}(u_i)$  and of the teacher model on augmented  $\mathcal{A}_g(u_i)$ , respectively,

$$p_i^t = f(\mathcal{A}_g(u_i); \theta_t), \quad (4)$$

$$p_i^s = f(\mathcal{T}(\mathcal{A}_g(u_i)); \theta_s). \quad (5)$$

Subsequently, the unlabeled loss is formulated as,

$$\mathcal{L}_u = \frac{1}{|\mathcal{B}_u|} \sum_{i=1}^{|\mathcal{B}_u|} \frac{1}{H \times W} \sum_{j=1}^{H \times W} \ell_{ce}(p_i^s(j), p_i^t(j)). \quad (6)$$

Different from recent SSS methods, our AugSeg follows a clean and simple two-branch teacher-student framework. We rely mainly on our augmentation strategy  $\mathcal{T}(\cdot)$  to produce prediction disagreement on the same input, which is also the key to semi-supervised learning. The augmentation  $\mathcal{T}$ , the core of AugSeg, consists of two kinds of augmentation in a cascade fashion, i.e.,  $\mathcal{T}(\cdot) = \mathcal{A}_r(\mathcal{A}_a(\cdot))$ , which are detailed in following sections.

### 3.2. Random Intensity-based Augmentations

In most existing semi-supervised learning studies, either in classification tasks or segmentation tasks, various auto augmentation techniques [9], especially the simplified Ran-

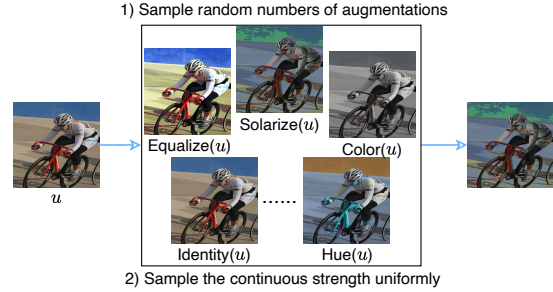


Figure 3. A visualization of random intensity-based augmentation.

domAug [10], have been widely adopted to perturb unlabeled samples. However, its different objective from semi-supervised learning limits its effectiveness. Specifically, the goal of various auto augmentations is to search for the optimal augmentation strategies for a specific downstream task. RandomAug further simplified this searching procedure in a finite discrete space. Whereas, the goal of data perturbation in semi-supervised learning is to generate two different views from the same image, where no specific optimal augmentation strategy is required. Besides, as discussed in [51], over-distorted augmentations will hurt the data distribution and degrade the SSS performance. To this end, we design a random intensity-based augmentation, denoted by  $\mathcal{A}_r$ , to perturb unlabeled data. As shown in Fig. 3, we

- sample the distorting degree uniformly in a continuous space instead of a finite discrete space.
- sample a random number of augmentations, bounded by a maximum value of  $k$ , from an augmentation pool instead of using a fixed number.
- remove strong intensity-based transformations like the Invert operations [51] in our augmentation pool. Our pool is directly simplified from the pool in RandomAug [11], as shown in Table 2.

In this way, our random intensity-based augmentation can enjoy better data diversity and adapt more to tsemi-supervised tasks. More importantly, different from [51], our highly random designs will not hurt the data distribution remarkably. Thus we can get rid of additional distribution-specific revisions [4] and extra filtering strategies [51].

### 3.3. Adaptive Label-aided CutMix

CutMix-related [13, 52] or copy-paste [18] augmentations have shown their effectiveness in supervised and semi-supervised segmentation tasks. Recent studies in SSS [17, 36, 47] apply the random copy-paste between unlabeled samples within a mini-batch and revise their pseudo-label accordingly. However, relying highly on the pseudo-labels

Method	Encoder	1/16 (92)	1/8 (183)	1/4 (366)	1/2 (732)	Full (1464)
Supervised	R50	44.03	52.26	61.65	66.72	72.94
PseudoSeg [58]	R50	54.89	61.88	64.85	70.42	71.00
PC <sup>2</sup> Seg [56]	R50	56.90	64.63	67.62	70.90	72.26
<b>AugSeg</b>	R50	<b>64.22</b>	<b>72.17</b>	<b>76.17</b>	<b>77.40</b>	<b>78.82</b>
Supervised	R101	43.92	59.10	65.88	70.87	74.97
CutMix-Seg [17]	R101	52.16	63.47	69.46	73.73	76.54
PseudoSeg [58]	R101	57.60	65.50	69.14	72.41	73.23
PC <sup>2</sup> Seg [56]	R101	57.00	66.28	69.78	73.05	74.15
CPS [7]	R101	64.07	67.42	71.71	75.88	-
PS-MT [36]	R101	65.80	69.58	76.57	78.42	80.01
ST++ [50]	R101	65.20	71.00	74.60	77.30	79.10
U <sup>2</sup> PL [47]	R101	67.98	69.15	73.66	76.16	79.49
<b>AugSeg</b>	R101	<b>71.09</b>	<b>75.45</b>	<b>78.80</b>	<b>80.33</b>	<b>81.36</b>

Table 3. Compared with the state-of-the-art methods on classic Pascal VOC 2012 val set under different partition protocols. ‘1/n’ means that ‘1/n’ data is used as labeled dataset, and the remaining images are used as unlabeled dataset.

may inevitably result in confirmation bias [2], especially for some difficult-to-train samples, or at the early training stages. Thus we tend to leverage the confident labeled samples to augment unlabeled data, so that labeled information can be fully exploited. However, mixing confident labeled information to unlabeled data is naturally beneficial but may under-utilize the unlabeled data. It is simply because some regions of unlabeled data are covered by regions from labeled samples and never utilized during the training. To this end, as shown in Figure 4, we design an adaptive label-aided augmentation that can make full use of labeled data to aid the training on unlabeled samples in an **instance-specific and confidence-adaptive** manner. In specific, we first estimate a confidence score,  $\rho_i$ , indicating the confidence level of the current model on its prediction on  $i$ -th unlabeled instance,

$$\rho_i = \frac{1}{H \times W} \sum_{j=1}^{H \times W} \max(p_i^t(j)) \left(1 - \frac{-\sum p_i^t(j) \log p_i^t(j)}{\log N}\right) \quad (7)$$

where we use the weighted average of the normalized prediction entropy on  $u_i$  to estimate the confidence score. Apparently, the score  $\rho$  is instance-specific and closely related to the generalization probability of the current model. We then use  $\rho_i$  as a triggering probability to randomly apply the mixing between labeled and unlabeled instances to obtain the mixing candidates  $\{u'_n\}$ ,

$$u'_n \leftarrow M_n \odot u_n + (\mathbf{1} - M_n) \odot x_n, \quad (8)$$

where  $M_n$  denotes the randomly generated region mask. After that, we apply the final mixing step between unlabeled instance  $\{u_m\}$  and the permuted mixing candidates  $\{u'_n\}$ ,

$$\mathcal{A}_a(u_m) \leftarrow M_m \odot u_m + (\mathbf{1} - M_m) \odot u'_n, \quad (9)$$

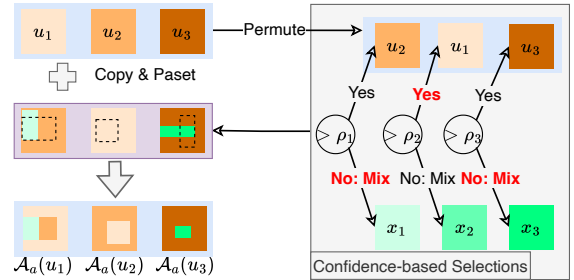


Figure 4. A visualization of adaptive label-aided CutMix augmentation in a mini-batch.  $x_i$  and  $u_i$  denote the labeled and unlabeled crops, respectively.  $\rho_i$  denote the confidence score for  $i$ -th unlabeled sample. The core idea of  $\mathcal{A}_a$  is that, these less confident unlabeled samples, with lower values of  $\rho_i$ , are more likely to be aided (mixed) by these confident labeled samples.

where  $M_m$ , similar to  $M_n$ , denotes the randomly generated binary region mask.

## 4. Experiments

In this section, we first describe our experimental setups and then compare our method with recent SOTAs on SSS benchmarks. A series of ablation studies are also conducted to verify the effectiveness and stability of AugSeg further.

### 4.1. Experimental setup

**Dataset.** We examine the efficacy of AugSeg on two widely used segmentation datasets, Pascal VOC 2012 [15] and Cityscapes [8]. Pascal VOC 2012 is a standard semantic segmentation benchmark with 21 semantic classes (including the background). The classic VOC 2012 includes 1,464 fine-labeled training images and 1,449 validating images. Following [7, 47, 50], we also include the

Method	ResNet-50			ResNet-101		
	1/16 (662)	1/8 (1323)	1/4 (2646)	1/16 (662)	1/8 (1323)	1/4 (2646)
Supervised	63.72	68.49	72.46	67.76	72.13	75.04
MT [46]	66.77	70.78	73.22	70.59	73.20	76.62
CCT [44]	65.22	70.87	73.43	67.94	73.00	76.17
GCT [29]	64.05	70.47	73.45	69.77	73.30	75.25
CPS [7]	68.21	73.20	74.24	72.18	75.83	77.55
CPS w/ CutMix [7]	71.98	73.67	74.90	74.48	76.44	77.68
ST++ [50]	72.60	74.40	75.40	74.50	76.30	76.60
PS-MT [36]	72.83	75.70	76.43	75.50	78.20	78.72
<b>AugSeg</b>	<b>74.66</b>	<b>75.99</b>	<b>77.16</b>	<b>77.01</b>	<b>77.31</b>	<b>78.82</b>
Supervised <sup>‡</sup>	67.66	71.91	74.53	70.63	75.02	76.47
U <sup>2</sup> PL <sup>‡*</sup> [47]	74.74	77.44	77.51	77.21	79.01	79.30
<b>AugSeg<sup>‡</sup></b>	<b>77.28</b>	<b>78.27</b>	<b>78.24</b>	<b>79.29</b>	<b>81.46</b>	<b>80.50</b>

Table 4. Comparison with the state-of-the-art on the PASCAL VOCAug val set under different partition protocols. The VOCAug trainset consists of 10,582 labeled samples in total. ‡ means the same split as U<sup>2</sup>PL, which prioritizes selecting high-quality labels from classic VOCs. Other methods use the same split as CPS. \* presents our reproduced results for U<sup>2</sup>PL [47] using ResNet-50.

additional coarsely-labeled 9,118 images from the Segmentation Boundary dataset (SBD) [22] as the training images, leading to 10,582 images in total. We adopt the same partition protocols in [7, 47] to evaluate our method on both `classic` and `blender` sets. Cityscapes consists of 19 semantic classes of urban scenes, including 2,975 training and 500 validating images with fine annotations.

**Training.** Following previous SSS studies [30, 36, 47], we adopt DeepLabV3+ [6] with ResNet [23] pretrained on ImageNet [12] as our segmentation backbone. Different from U<sup>2</sup>PL [47], we use an output stride of 16 by default (instead of using 8). We use an SGD optimizer with a momentum of 0.9 and a polynomial learning-rate decay with an initial value of 0.01 to train the student model. Referring to [7, 47], we use the crop size of  $512 \times 512$  with a training epoch of 80 and the crop size of  $800 \times 800$  with a training epoch of 240 on VOC2012 and Cityscapes, respectively. A batch size of 16 and the sync-BN are adopted for both datasets.

**Evaluations.** We use the mean of intersection-over-union to evaluate the segmentation performance for all runs, using ResNet-50 and ResNet-101 as the encoder separately. Following CPS [7] and U<sup>2</sup>PL [47], we also adopt the sliding evaluation to examine the performance on validation images of Cityscapes with a resolution of  $1024 \times 2048$ .

## 4.2. Comparison with SOTAs

We demonstrate the superiority of AugSeg by comparing it with current SOTAs on both datasets under different partition protocols. Since U<sup>2</sup>PL prioritizes selecting high-quality labels from classic VOCs for testing on `blender` VOC<sup>2</sup>, we reproduce the supervised baseline and its perfor-

mance on ResNet-50 for fair comparisons.

**Pascal VOC 2012.** In Table 3 and Table 4, we compare our AugSeg with recent SOTAs on classic and `blender` VOC, respectively. It can be clearly seen from Table 3 that, despite its simplicity, AugSeg can consistently outperform current SOTAs by a large margin, *e.g.* obtaining a 4.45% performance gain on 1/8 split using R101 as the encoder. Note that AugSeg can even achieve higher performance of 71.09% with only 92 labels than the previous SOTA performance of 71.00% with 183 labels. We can also observe that the performance gains become more noticeable and obvious when using ResNet-50 as the encoder and when fewer labels, *e.g.* 92 and 182 labels, are available. Even though the performance gap between different SSS methods is decreasing as more labeled data is involved, our method can still improve the previous SOTA by 1.91% and 1.35% with 1/2 and full fine annotations, respectively.

As shown in Table 4, under 1/16, 1/8 and 1/4 partition protocols with the same split as U<sup>2</sup>PL, our AugSeg obtains new SOTA performance of 79.29%, 81.46%, and 80.50% based on ResNet-101, which obtains around 2% performance improvements again previous SOTA. It is noteworthy that AugSeg obtains a higher mIoU with 1323 labels than that with 2646. It is simply because the 2646 split involves more noisy (coarsely-annotated) labels than 1323 split (there are 1464 fine annotations in total). As discussed in Sec. 3.3, AugSeg can make full use of labeled data to stabilize the training on unlabeled data. Thus the noisy labeled information will degrade the performance of AugSeg. This is also why the superiority of AugSeg is more noticeable in Table 3 than in Table 4 with the same split as CPS [7].

**Cityscapes.** In Table 5, we evaluate our method on

<sup>2</sup><https://github.com/Haochen-Wang409/U2PL/issues/3>

Method	ResNet-50				ResNet-101			
	1/16(186)	1/8(372)	1/4(744)	1/2(1488)	1/16(186)	1/8(372)	1/4(744)	1/2(1488)
Supervised	63.34	68.73	74.14	76.62	64.77	71.64	75.24	78.03
MT [46]	66.14	72.03	74.47	77.43	68.08	73.71	76.53	78.59
CCT [44]	66.35	72.46	75.68	76.78	69.64	74.48	76.35	78.29
GCT [29]	65.81	71.33	75.30	77.09	66.90	72.96	76.45	78.58
CPS [7]	69.79	74.39	76.85	78.64	70.50	75.71	77.41	80.08
CPS * [7]	-	-	-	-	69.78	74.31	74.58	76.81
PS-MT† [36]	-	75.76	76.92	77.64	-	76.89	77.60	79.09
U <sup>2</sup> PL [47]	69.03	73.02	76.31	78.64	70.30	74.37	76.47	79.05
<b>AugSeg</b>	<b>73.73</b>	<b>76.49</b>	<b>78.76</b>	<b>79.33</b>	<b>75.22</b>	<b>77.82</b>	<b>79.56</b>	<b>80.43</b>

Table 5. Comparison with state-of-the-art on Cityscapes val set under different partition protocols. Cityscapes (Citys) includes 2, 975 samples in total. \* means the reproduced results in U<sup>2</sup>PL [47]. All the results are reported by the sliding evaluations. † means PS-MT [36] runs more epochs (320, 450, 550 epochs on 1/8, 1/4, 1/2 splits, respectively) than ours (240 epochs for all splits.)

more challenging Cityscapes, using ResNet-50 and ResNet-101 as the encoder, respectively. We can easily see that AugSeg can readily outperform other SSS methods, especially with scarce labels. Though AugSeg is embarrassingly simpler than the recent SOTA U<sup>2</sup>PL in terms of the training procedure and encoded feature size, AugSeg can improve U<sup>2</sup>PL by 4.92%, 3.45%, 3.09%, and 1.38%, using ResNet-101 as the encoder, under 1/16, 1/8, 1/4, and 1/2 partition protocols, respectively. Not relying on advanced unsupervised techniques or multiple trainable models, AugSeg can consistently achieve the best performance on SSS benchmarks. Such impressive performance improvement further demonstrates the effectiveness and importance of our claim that various data augmentation should be simplified and adjusted to better adapt to semi-supervised learning.

In addition, we highlight the importance of labeled samples in semi-supervised learning, in terms of the quantity and quality. First, regardless of different semi-supervised approaches, we can see from Tables 3, 4 and 5 that providing more labeled samples can easily boost the semi-supervised performance. Second, comparing the performance on classic and blended VOCs, we observe that the quality of labeled samples is always crucial. For example, our AugSeg can achieve a high performance of 78.80% using 366 high-quality labels but require 2646 labels from the blender dataset to obtain a comparable mIOU of 78.82%.

### 4.3. Ablations studies

In this section, we conduct a series of ablations studies on Pascal VOC 2012 with 366 (1/4) labels and Cityscape with 744 (1/4) labels using ResNet-50 as the encoder.

**Effectiveness of different components of AugSeg.** We first investigate the effectiveness of each component of AugSeg in Table 6. Our simplified augmentations  $\mathcal{A}_r$  and  $\mathcal{A}_a$  can effectively improve the SSS performance, obtain-

AugSeg			mIoU	
MT	$\mathcal{A}_r$	$\mathcal{A}_a$	VOC (366)	Citys (744)
			61.65 (supervised)	74.14 (supervised)
✓			69.06 (7.41↑)	75.96 (1.82↑)
✓	✓		72.41 (10.76↑)	77.29 (3.15↑)
✓		✓	74.33 (12.68↑)	77.44 (3.30↑)
✓	✓	✓	<b>76.17 (14.52↑)</b>	<b>78.76(4.62↑)</b>

Table 6. Ablation studies on our AugSeg. “MT” means the standard mean-teacher semi-supervised framework.  $\mathcal{A}_r$  and  $\mathcal{A}_a$  represent the two main augmentation strategies, the random intensity-based and adaptive label-aided augmentations, respectively. Improvements over the supervised baseline are highlighted in blue.




Figure 5. Ablation studies on different designs of AugSeg, where “LargePool” refers to the augmentation pool in [51].

ing 10.76% and 12.68% improvements against the supervised baseline, 3.35% and 5.27% improvements against the plain mean-teacher, on VOC 2012 with 366 labels, respectively. Using adaptive label-aided augmentation  $\mathcal{A}_a$  can obtain better performance than using  $\mathcal{A}_r$  individually. Integrating both augmentations can further improve each individual component and achieve the best performance.

**Impact of different augmentation designs.** As shown in Figure 5, we test the impact of our simplified and adap-

$\lambda_u$	0.0	0.5	1.0	1.5	2.0
VOC (366)	61.65	75.21	76.17	75.95	<b>77.05</b>
Citys (744)	74.14	77.02	78.76	<b>78.99</b>	78.68

Table 7. Ablations on the loss weight  $\lambda_u$ , set as 1.0 by default.

$k$	0	1	2	3	4
VOC (366)	74.38	75.50	76.10	76.17	<b>76.32</b>
Citys (186)	71.26	72.10	73.42	<b>73.73</b>	73.03
Citys (744)	77.44	78.34	78.11	<b>78.76</b>	78.48

Table 8. Ablations on the maximum number of selected intensity-based augmentations, using R50 as the encoder.  $k = 3$  by default.

tive designs on intensity-based and cutmix-based augmentations. We adopt  $\mathcal{A}_a$  to stabilize the training on unlabeled data adaptively, which benefits the SSS training apparently. Following SimpleBaseline [51], we also examine our random intensity-based designs with a larger augmentation pool. Consequently, although more augmentation selections can improve individual performance, incorporating both augmentations with more strong augmentation selections can degrade the performance, resulting from the discussed over-distortion issues in [51]. Similar observations can be found from Figure 6 that using a fixed strategy to select more augmentations can harm the SSS performance. Our highly random and simplified designs can naturally alleviate this issue without introducing extra operations like Distribution-specific BN [4].

**Impact of hyper-parameters.** We also examine the segmentation performance of AugSeg with different maximum numbers of selected augmentations,  $k$  and different values of unsupervised loss weight  $\lambda_u$ , in Table 8 and Table 7, respectively.  $\lambda_u = 0$  means no training on unlabeled data while  $k = 0$  means not applying the rand intensity-based augmentation. We can observe that thanks to our random design, a larger  $k$  will not degrade but enhance the SSS performance, compared to the results with fixed selecting in Figure 6. It can also be seen from both tables that no specific parameters can consistently outperform others on both datasets. To keep our method simple and consistent, we set  $k = 3$  and  $\lambda_u = 1.0$  for all runs by default.

**Qualitative Results.** Figure 7 shows some qualitative results on the Pascal VOC 2012 dataset. The supervised baseline obtains the worse segmentation results, *e.g.*, not capable of differentiating the train and bus. Using more unlabeled data in the plain MT can enhance the model’s capability to separate confusing classes. Among the comparisons, some challenging small-sized objects, such as the wheels, grass, or humans in a large background, can only be effectively identified by our AugSeg, which further demon-




Figure 6. Impact of different selecting strategies in intensity-based augmentations with different numbers of selected operations.



Figure 7. Qualitative results on Pascal VOC 2012 with 366 labels and ResNet-50 as the encoder. Columns from left to right denote the original images, the ground-truth, the supervised results, the plain MT results, and our AugSeg results, respectively.

strates the effectiveness of AugSeg. Though our method obtains the SOTA SSS performance, AugSeg is **limited** at identifying some hard-to-segment objects, *e.g.*, cars in the advertisement. We believe there is still great potential to further improve the SSS performance on top of our AugSeg.

## 5. Conclusion

In this paper, we propose AugSeg, a simple-yet-effective approach to semi-supervised semantic segmentation. Unlike recent SSS studies that tend to combine increasingly complicated mechanisms, AugSeg follows a standard two-branch teacher-student framework to train models on labeled and unlabeled data jointly. The key to AugSeg lies in the simplification and revisions of two existing augmentation models, *i.e.*, the random intensity-based and adaptive label-aided CutMix-based augmentations. Without any additional complicated designs, AugSeg readily obtains new SOTA performance on popular SSS benchmarks under different partition protocols. We hope our AugSeg can serve as a strong baseline for future SSS studies.



## References

- [1] Inigo Alonso, Alberto Sabater, David Ferstl, Luis Monteseano, and Ana C Murillo. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In *ICCV*, 2021. 1
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks*. IEEE, 2020. 5
- [3] Eric Arazo, Diego Ortego, Paul Albert, Noel E O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *IJCNN*, 2020. 2
- [4] Woong-Gi Chang, Tackgeun You, Seonguk Seo, Suha Kwak, and Bohyung Han. Domain-specific batch normalization for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 7354–7362, 2019. 3, 4, 8
- [5] 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.*, 2017. 1
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 1, 6
- [7] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *CVPR*, 2021. 1, 2, 5, 6, 7
- [8] 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 *CVPR*, 2016. 5
- [9] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018. 1, 3, 4
- [10] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 1, 2, 3, 4
- [11] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPRW*, 2020. 3, 4
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [13] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv:1708.04552*, 2017. 4
- [14] Yue Duan, Zhen Zhao, Lei Qi, Lei Wang, Luping Zhou, Yinghuan Shi, and Yang Gao. Mutexmatch: semi-supervised learning with mutex-based consistency regularization. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 2
- [15] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *IJCV*, 2015. 5
- [16] Zhengyang Feng, Qianyu Zhou, Qiqi Gu, Xin Tan, Guangliang Cheng, Xuequan Lu, Jianping Shi, and Lizhuang Ma. Dmt: Dynamic mutual training for semi-supervised learning. *arXiv:2004.08514*, 2020. 1
- [17] Geoff French, Timo Aila, Samuli Laine, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, high-dimensional perturbations. In *BMVC*, 2020. 1, 2, 3, 4, 5
- [18] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, 2021. 2, 4
- [19] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *NeurIPS*, 2005. 2
- [20] Dayan Guan, Jiaying Huang, Aoran Xiao, and Shijian Lu. Unbiased subclass regularization for semi-supervised semantic segmentation. In *CVPR*, 2022. 2, 3
- [21] Guan Gui, Zhen Zhao, Lei Qi, Luping Zhou, Lei Wang, and Yinghuan Shi. Improving barely supervised learning by discriminating unlabeled samples with super-class. In *Advances in Neural Information Processing Systems*, 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. 6
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [24] Ruifei He, Jihan Yang, and Xiaojuan Qi. Re-distributing biased pseudo labels for semi-supervised semantic segmentation: A baseline investigation. In *ICCV*, 2021. 2
- [25] Hanzhe Hu, Fangyun Wei, Han Hu, Qiwei Ye, Jinshi Cui, and Liwei Wang. Semi-supervised semantic segmentation via adaptive equalization learning. In *NeurIPS*, 2021. 2
- [26] Lin Hu, Jiabin Li, Xingchen Peng, Jianghong Xiao, Bo Zhan, Chen Zu, Xi Wu, Jiliu Zhou, and Yan Wang. Semi-supervised npc segmentation with uncertainty and attention guided consistency. *Knowledge-Based Systems*, 239:108021, 2022. 2
- [27] Wei-Chih Hung, Yi-Hsuan Tsai, Yan-Ting Liou, Yen-Yu Lin, and Ming-Hsuan Yang. Adversarial learning for semi-supervised semantic segmentation. In *BMVC*, 2018. 1, 3
- [28] Mostafa S Ibrahim, Arash Vahdat, Mani Ranjbar, and William G Macready. Semi-supervised semantic image segmentation with self-correcting networks. In *CVPR*, 2020. 1
- [29] Zhanghan Ke, Kaican Li Di Qiu, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. In *ECCV*, 2020. 1, 6, 7
- [30] Donghyeon Kwon and Suha Kwak. Semi-supervised semantic segmentation with error localization network. In *CVPR*, 2022. 1, 2, 3, 6
- [31] Xin Lai, Zhuotao Tian, Li Jiang, Shu Liu, Hengshuang Zhao, Liwei Wang, and Jiaya Jia. Semi-supervised semantic seg-

- mentation with directional context-aware consistency. In *CVPR*, 2021. 2
- [32] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 2
- [33] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop*, 2013. 2
- [34] Shikun Liu, Shuaifeng Zhi, Edward Johns, and Andrew J Davison. Bootstrapping semantic segmentation with regional contrast. In *ICLR*, 2022. 2
- [35] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Zhaoyu Wang, Li Mian, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *arXiv preprint arXiv:2006.08218*, 1(2), 2020. 3
- [36] Yuyuan Liu, Yu Tian, Yuanhong Chen, Fengbei Liu, Vasileios Belagiannis, and Gustavo Carneiro. Perturbed and strict mean teachers for semi-supervised semantic segmentation. In *CVPR*, 2022. 1, 2, 3, 4, 5, 6, 7
- [37] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1
- [38] Robert Mendel, Luis Antonio de Souza, David Rauber, João Paulo Papa, and Christoph Palm. Semi-supervised segmentation based on error-correcting supervision. In *ECCV*, 2020. 1, 2, 3
- [39] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high-and low-level consistency. *IEEE Transactions on Pattern Analysis and Machine Intelligence.*, 2019. 1
- [40] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. 3
- [41] Samuel G. Müller and Frank Hutter. Trivialaugument: Tuning-free yet state-of-the-art data augmentation. In *ICCV*, 2021. 3
- [42] Avital Oliver, Augustus Odena, Colin Raffel, Ekin D Cubuk, and Ian J Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. *arXiv preprint arXiv:1804.09170*, 2018. 1, 2
- [43] Yassine Ouali, Céline Hudelot, and Myriam Tami. An overview of deep semi-supervised learning. *arXiv preprint arXiv:2006.05278*, 2020. 1, 3
- [44] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *CVPR*, 2020. 1, 2, 3, 6, 7
- [45] 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 *NeurIPS*, 2020. 2
- [46] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 1, 2, 3, 6, 7
- [47] 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 *CVPR*, 2022. 1, 2, 3, 4, 5, 6, 7
- [48] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. In *NeurIPS*, 2020. 2, 3
- [49] 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
- [50] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. St++: Make self-training work better for semi-supervised semantic segmentation. In *CVPR*, 2022. 1, 2, 5, 6
- [51] Jianlong Yuan, Yifan Liu, Chunhua Shen, Zhibin Wang, and Hao Li. A simple baseline for semi-supervised semantic segmentation with strong data augmentation. In *ICCV*, 2021. 1, 2, 3, 4, 7, 8
- [52] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 1, 4
- [53] 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. 1
- [54] Zhen Zhao, Luping Zhou, Yue Duan, Lei Wang, Lei Qi, and Yinghuan Shi. Dc-ssl: Addressing mismatched class distribution in semi-supervised learning. In *CVPR*, 2022. 2
- [55] Zhen Zhao, Luping Zhou, Lei Wang, Yinghuan Shi, and Yang Gao. Lassl: Label-guided self-training for semi-supervised learning. In *AAAI*, 2022. 2
- [56] Yuanyi Zhong, Bodi Yuan, Hong Wu, Zhiqiang Yuan, Jian Peng, and Yu-Xiong Wang. Pixel contrastive-consistent semi-supervised semantic segmentation. In *ICCV*, 2021. 1, 2, 5
- [57] 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 *ICCV*, 2021. 1, 2
- [58] 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 *ICLR*, 2021. 2, 5