

Revisiting Weak-to-Strong Consistency in Semi-Supervised Semantic Segmentation

Lihe Yang¹ Lei Qi² Litong Feng³ Wayne Zhang³ Yinghuan Shi^{1*}
¹Nanjing University ²Southeast University ³SenseTime Research

<https://github.com/LiheYoung/UniMatch>

Abstract

In this work, we revisit the weak-to-strong consistency framework, popularized by FixMatch from semi-supervised classification, where the prediction of a weakly perturbed image serves as supervision for its strongly perturbed version. Intriguingly, we observe that such a simple pipeline already achieves competitive results against recent advanced works, when transferred to our segmentation scenario. Its success heavily relies on the manual design of strong data augmentations, however, which may be limited and inadequate to explore a broader perturbation space. Motivated by this, we propose an auxiliary feature perturbation stream as a supplement, leading to an expanded perturbation space. On the other, to sufficiently probe original image-level augmentations, we present a dual-stream perturbation technique, enabling two strong views to be simultaneously guided by a common weak view. Consequently, our overall Unified Dual-Stream Perturbations approach (UniMatch) surpasses all existing methods significantly across all evaluation protocols on the Pascal, Cityscapes, and COCO benchmarks. Its superiority is also demonstrated in remote sensing interpretation and medical image analysis. We hope our reproduced FixMatch and our results can inspire more future works.

1. Introduction

Semantic segmentation aims to provide pixel-level predictions to images, which can be deemed as a dense classification task and is fundamental to real-world applications, e.g., autonomous driving. Nevertheless, conventional fully-supervised scenario [43, 73, 77] is extremely hungry for delicately labeled images by human annotators, greatly hindering its broad application to some fields where it is costly and even infeasible to annotate abundant images. Therefore, semi-supervised semantic segmentation [56] has been proposed and is attracting increasing attention. Generally, it wishes to alleviate the labor-intensive process via leverag-

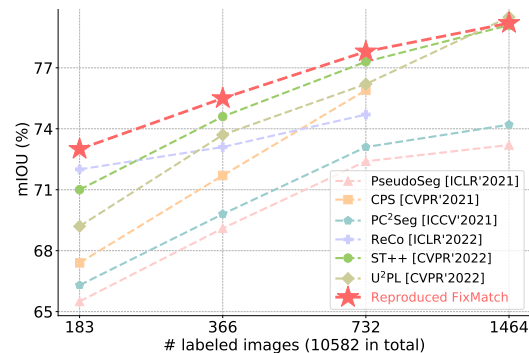


Figure 1. Comparison between state-of-the-art methods and our reproduced FixMatch [55] on the Pascal dataset.

ing a large quantity of unlabeled images, accompanied by a handful of manually labeled images.

Following closely the research line of semi-supervised learning (SSL), advanced methods in semi-supervised semantic segmentation have evolved from GANs-based adversarial training paradigm [21, 47, 56] into the widely adopted consistency regularization framework [13, 19, 28, 29, 49, 61, 81] and reborn self-training pipeline [23, 27, 68, 70]. In this work, we focus on the weak-to-strong consistency regularization framework, which is popularized by FixMatch [55] from the field of semi-supervised classification, and then impacts many other relevant tasks [42, 45, 57, 62, 66, 67]. The weak-to-strong approach supervises a strongly perturbed unlabeled image x^s with the prediction yielded from its corresponding weakly perturbed version x^w , as illustrated in Figure 2a. Intuitively, its success lies in that the model is more likely to produce high-quality prediction on x^w , while x^s is more effective for our model to learn, since the strong perturbations introduce additional information as well as mitigate confirmation bias [2]. We surprisingly notice that, so long as coupled with appropriate strong perturbations, FixMatch can indeed still exhibit powerful generalization capability in our scenario, obtaining superior results over state-of-the-art (SOTA) methods, as compared in Figure 1. Thus, we select this simple yet effective framework as our baseline.

Through investigation of image-level strong perturbations,

*Corresponding author.

Method	# labeled images (10582 in total)				
	92	183	366	732	1464
w/o any SP	39.5	52.7	65.5	69.2	74.6
w/ CutMix	<u>56.7</u>	<u>67.9</u>	<u>71.9</u>	<u>75.1</u>	<u>78.3</u>
w/ whole SP	63.9	73.0	75.5	77.8	79.2

Table 1. The importance of image-level strong perturbations (SP) to FixMatch on the Pascal dataset. *w/o any SP*: directly utilize hard label of x^w to supervise its logits. *w/ CutMix*: only use CutMix [71] as a perturbation. *w/ whole SP*: strong perturbations contain color transformations from ST++ [68], together with CutMix.

we observe that they play an indispensable role in making the FixMatch a rather strong competitor in semi-supervised semantic segmentation. As demonstrated in Table 1, the performance gap between whether to adopt perturbations is extremely huge. Greatly inspired by these clues, we hope to inherit the spirit of strong perturbations from FixMatch, but also further strengthen them from two different perspectives and directions, namely *expanding a broader perturbation space*, and *sufficiently harvesting original perturbations*. Each of these two perspectives is detailed in the following two paragraphs respectively.

Image-level perturbations, *e.g.*, color jitter and CutMix [71], include heuristic biases, which actually introduce additional prior information into the bootstrapping paradigm of FixMatch, so as to capture the merits of consistency regularization. In case not equipped with these perturbations, FixMatch will be degenerated to a naive online self-training pipeline, producing much worse results. Despite its effectiveness, these perturbations are totally constrained at the image level, hindering the model to explore a broader perturbation space and to maintain consistency at diverse levels. To this end, in order to expand original perturbation space, we design a unified perturbation framework for both raw images and extracted features. Concretely, on raw images, similar to FixMatch, pre-defined image-level strong perturbations are applied, while for extracted features of weakly perturbed images, an embarrassingly simple channel dropout is inserted. In this way, our model pursues the equivalence of predictions on unlabeled images at both the image and embedding level. These two perturbation levels can be complementary to each other. Distinguished from [33, 41], we separate different levels of perturbations into independent streams to avoid a single stream being excessively hard to learn.

On the other hand, current FixMatch framework merely utilizes a single strong view of each unlabeled image in a mini-batch, which is insufficient to fully exploit the manually pre-defined perturbation space. Considering this, we present a simple yet highly effective improvement to the input, where dual independent strong views are randomly sampled from the perturbation pool. They are then fed into the student model in parallel, and simultaneously supervised by their

shared weak view. Such a minor modification even easily turns the FixMatch baseline into a stronger SOTA framework by itself. Intuitively, we conjecture that enforcing two strong views to be close to a common weak view can be regarded as minimizing the distance between these strong views. Hence, it shares the spirits and merits of contrastive learning [11, 25], which can learn more discriminative representations and is proved to be particularly beneficial to our current task [40, 61]. We conduct comprehensive studies on the effectiveness of each proposed component. Our contributions can be summarized in four folds:

- We notice that, coupled with appropriate image-level strong perturbations, FixMatch is still a powerful framework when transferred to the semantic segmentation scenario. A plainly reproduced FixMatch outperforms almost all existing methods in our current task.
- Built upon FixMatch, we propose a unified perturbation framework that unifies image-level and feature-level perturbations in independent streams, to exploit a broader perturbation space.
- We design a dual-stream perturbation strategy to fully probe pre-defined image-level perturbation space, as well as to harvest the merits of contrastive learning for discriminative representations.
- Our framework that integrates above two components, surpasses existing methods remarkably across all evaluation protocols on the Pascal, Cityscapes, and COCO. Notably, it also exhibits strong superiority in medical image analysis and remote sensing interpretation.

2. Related Work

Semi-supervised learning (SSL). The core issue in SSL lies in how to design reasonable and effective supervision signals for unlabeled data. Two main branches of methodology are proposed to tackle the issue, namely entropy minimization [22, 37, 51, 53, 64, 80] and consistency regularization [5, 6, 20, 30, 36, 38, 48, 54, 58, 63]. Entropy minimization, popularized by self-training [37], works in a straightforward way via assigning pseudo labels to unlabeled data and then combining them with manually labeled data for further re-training. For another thing, consistency regularization holds the assumption that prediction of an unlabeled example should be invariant to different forms of perturbations. Among them, FixMatch [55] proposes to inject strong perturbations to unlabeled images and supervise training process with predictions from weakly perturbed ones to subsume the merits of both methodologies. Recently, FlexMatch [72] and FreeMatch [60] consider learning status of different classes and then filter low-confidence labels with class-wise thresholds. Our method inherits from FixMatch, however, we investigate a more challenging and labor-intensive setting. More importantly, we demonstrate the significance

of image-level strong perturbations, thereby managing to expand original perturbation space and take full advantage of pre-defined perturbations.

Semi-supervised semantic segmentation. Earlier works [47, 56] incorporate the GANs [21] as an auxiliary supervision for unlabeled images via discriminating pseudo labels from manual labels. Motivated by the rapid progress in SSL, recent methods [1, 18, 32, 34, 40, 41, 46, 49, 74, 75, 78, 79, 81] strive for simpler training paradigms from the perspective of consistency regularization and entropy minimization. During this trend, French *et al.* [19] disclose Cutout [16] and CutMix [71] are critical to success of consistency regularization in segmentation. AEL [28] then designs an adaptive CutMix and sampling strategy to enhance the learning on underperforming classes. Inspired by contrastive learning, Lai *et al.* [35] propose to enforce predictions of the shared patch under different contextual crops to be same. And U²PL [61] treats uncertain pixels as reliable negative samples to contrast against corresponding positive samples. Similar to the core spirit of co-training [7, 52, 76], CPS [13] introduces dual models to supervise each other.

Other works from the research line of entropy minimization utilize a self-training pipeline to assign pseudo masks for unlabeled images in an offline manner. From this perspective, Yuan *et al.* [70] claim excessive perturbations on unlabeled images are catastrophic to clean data distribution, and thus propose a separate batch normalization for these images. Concurrently, ST++ [68] points out that appropriate strong data perturbations are indeed extremely helpful to self-training. Moreover, to tackle the class bias issue encountered in pseudo labeling, He *et al.* [27] align class distributions between manual labels and pseudo labels. And USRN [23] clusters balanced subclass distributions as a regularization to alleviate the imbalance of pre-defined classes.

To pursue elegance and efficacy, we adopt the weak-to-strong consistency regularization framework from FixMatch [55]. Our end-to-end baseline can be deemed as an improvement of [19], or a simplification of [81]. For instance, it strengthens image-level strong perturbations in [19] with color transformations from [68], and discards the calibration fusion module in [81]. With this neat but competitive baseline, we further probe a broader perturbation space, and fully exploit original image-level perturbations as well.

3. Method

Algorithms in semi-supervised semantic segmentation aim to fully explore unlabeled images $\mathcal{D}^u = \{x_i^u\}$ with limited amount of annotations from labeled images $\mathcal{D}^l = \{(x_i^l, y_i^l)\}$. Our method is based on FixMatch [55], so we first briefly review its core idea (Sec. 3.1). Following this, we introduce the two proposed components in detail, namely unified perturbations (Sec. 3.2), and dual-stream perturbations (Sec. 3.3). Finally, we summarize our overall Unified

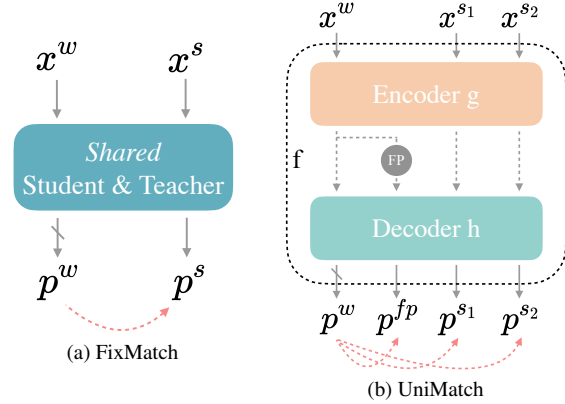


Figure 2. (a) The FixMatch baseline. (b) Our proposed unified dual-stream perturbations method (UniMatch). The FP denotes feature perturbation, and the dashed curves represent supervision.

Dual-Stream Perturbations method (UniMatch) (Sec. 3.4).

3.1. Preliminaries

As aforementioned, FixMatch utilizes a weak-to-strong consistency regularization to leverage unlabeled data. Concretely, each unlabeled image x^u is simultaneously perturbed by two operators, *i.e.*, weak perturbation \mathcal{A}^w such as cropping, and strong perturbation \mathcal{A}^s such as color jitter. Then, the overall objective function is a combination of supervised loss \mathcal{L}_s and unsupervised loss \mathcal{L}_u as:

$$\mathcal{L} = \frac{1}{2}(\mathcal{L}_s + \mathcal{L}_u). \quad (1)$$

Typically, the supervised term \mathcal{L}_s is the cross-entropy loss between model predictions and groundtruth labels. And the unsupervised loss \mathcal{L}_u regularizes prediction of the sample under strong perturbations to be the same as that under weak perturbations, which can be formulated as:

$$\mathcal{L}_u = \frac{1}{B_u} \sum \mathbb{1}(\max(p^w) \geq \tau) H(p^w, p^s), \quad (2)$$

where B_u is the batch size for unlabeled data and τ is a pre-defined confidence threshold to filter noisy labels. H minimizes the entropy between two probability distributions:

$$p^w = \hat{F}(\mathcal{A}^w(x^u)); \quad p^s = F(\mathcal{A}^s(\mathcal{A}^w(x^u))), \quad (3)$$

where the teacher model \hat{F} produces pseudo labels on weakly perturbed images, while the student F leverages strongly perturbed images for model optimization. In this work, we set \hat{F} exactly the same as F for simplicity, following FixMatch.

3.2. Unified Perturbations for Images and Features

Apart from semi-supervised classification, the methodology in FixMatch has swept across a wide range of research topics and achieved booming success, such as semantic segmentation [19, 28, 81], object detection [42, 57, 66], unsupervised domain adaptation [45], and action recognition [62, 67].




Figure 3. The two sub-frameworks, *i.e.*, UniPerb and DusPerb, that are integrated into our final UniMatch framework. The FP denotes feature perturbation, which is a simple Dropout in this work.

Despite its popularity, its efficacy actually heavily depends on delicately designed strong perturbations from researchers, whose optimal combinations and hyper-parameters are time-consuming to acquire. Besides, in some cases such as medical image analysis and remote sensing interpretation, it may require domain-specific knowledge to figure out promising ones. More importantly, they are completely constrained at the image level, hindering the student model to maintain multi-level consistency against more diverse perturbations.

To this end, in order to construct a broader perturbation space, built on top of FixMatch, we propose to inject perturbations on features of the *weakly* perturbed image x^w . We choose to separate different levels of perturbations into multiple independent feedforward streams, enabling the student to achieve targeted consistency in each stream more directly. Formally, a segmentation model f can be decomposed into an encoder g and a decoder h . In addition to acquired p^w and p^s in FixMatch, we also obtain p^{fp} from an auxiliary feature perturbation stream by:

$$e^w = g(x^w), \quad (4)$$

$$p^{fp} = h(\mathcal{P}(e^w)), \quad (5)$$

where e^w is extracted features of x^w , and \mathcal{P} denotes feature perturbations, *e.g.*, dropout or adding uniform noise.

Overall, as exhibited in Figure 3a, three feedforward streams are maintained for each unlabeled mini-batch, which are (i) the simplest stream: $x^w \rightarrow f \rightarrow p^w$, (ii) image-level strong perturbation stream: $x^s \rightarrow f \rightarrow p^s$, and (iii) our introduced feature perturbation stream: $x^w \rightarrow g \rightarrow \mathcal{P} \rightarrow h \rightarrow p^{fp}$. In this way, the student model is enforced to be consistent to unified perturbations at both image and feature level. We name it as UniPerb for convenience. The unsupervised loss \mathcal{L}_u is formulated as:

$$\mathcal{L}_u = \frac{1}{B_u} \sum \mathbb{1}(\max(p^w) \geq \tau) (\mathbb{H}(p^w, p^s) + \mathbb{H}(p^w, p^{fp})). \quad (6)$$

It should be noted that, we do not aim at proposing a novel feature perturbation approach in this work. Actually, an

Algorithm 1 Pseudocode of UniMatch in a PyTorch-like style.

```
# f: network, composed of an encoder g and a decoder h
# aug_w/aug_s: weak/strong image-level perturbations

for x in loader_u:
    # one weak view and two strong views as input
    x_w = aug_w(x)
    x_s1, x_s2 = aug_s(x_w), aug_s(x_w)

    # feature of weakly perturbed image
    feat_w = g(x_w)
    # perturbed feature
    feat_fp = nn.Dropout2d(0.5)(feat_w)

    # four predictions from four forward streams
    p_w, p_fp = h(torch.cat((feat_w, feat_fp))).chunk(2)
    p_s1, p_s2 = f(torch.cat((x_s1, x_s2))).chunk(2)

    # hard (one-hot) pseudo mask
    mask_w = p_w.argmax(dim=1).detach()

    # loss from image- and feature-level perturbation
    criterion = nn.CrossEntropyLoss()
    p_s = torch.cat((p_s1, p_s2))
    loss_s = criterion(p_s, mask_w.repeat(2, 1, 1))
    loss_fp = criterion(p_fp, mask_w)

    # final unsupervised loss
    loss_u = (loss_s + loss_fp) / 2.0
```

embarrassingly simple channel dropout (`nn.Dropout2d` in PyTorch) is well-performed enough. Furthermore, distinguished from recent work [41] that mixes three levels of perturbations into a single stream, we highlight the necessity of separating perturbations of different properties into independent streams, which is evidenced in our ablation studies. We believe that, image-level perturbations can be well complemented by feature-level perturbations.

3.3. Dual-Stream Perturbations

Motivated by the tremendous advantages of image-level strong perturbations, we wish to fully explore them. We are inspired by recent progress in self-supervised learning and semi-supervised classification, that constructing multiple views for unlabeled data as inputs can better leverage the perturbations. For instance, SwAV [8] proposes a novel technique called `multi-crop`, enforcing the local-to-global consistency among a bag of views of different resolutions. Likewise, ReMixMatch [5] produces multiple strongly augmented versions for the model to learn.

Therefore, we wonder whether such a simple idea can also benefit our semi-supervised semantic segmentation. We make a straightforward attempt that, rather than feeding a single p^s into the model, we independently yield dual-stream perturbations (x^{s1}, x^{s2}) from x^w by strong perturbation pool \mathcal{A}^s . Since \mathcal{A}^s is pre-defined but non-deterministic, x^{s1} and x^{s2} are not equal. This dual-stream perturbation framework (DusPerb) is displayed in Figure 3b.

Intriguingly, such a minor modification brings consistent and substantial improvements over original FixMatch under all partition protocols in our segmentation scenario, establishing new state-of-the-art results. It is validated in our ablation studies that, the performance gain is non-trivial, not

Pascal	92	183	366	732	1464
SupBaseline	44.0	52.3	61.7	66.7	72.9
PC ² Seg [78] <small>[ICCV'21]</small>	56.9	64.6	67.6	70.9	72.3
UniMatch RN-50	71.9	72.5	76.0	77.4	78.7
SupBaseline	45.1	55.3	64.8	69.7	73.5
CPS [13] <small>[CVPR'21]</small>	64.1	67.4	71.7	75.9	-
ST++ [68] <small>[CVPR'22]</small>	65.2	71.0	74.6	77.3	79.1
U ² PL [61] <small>[CVPR'22]</small>	68.0	69.2	73.7	76.2	79.5
PS-MT [41] <small>[CVPR'22]</small>	65.8	69.6	76.6	78.4	80.0
PCR [65] <small>[NeurIPS'22]</small>	70.1	74.7	77.2	78.5	80.7
UniMatch RN-101	75.2	77.2	78.8	79.9	81.2

Table 2. Comparison with SOTAs on the **Pascal**. Labeled images are from the *original high-quality* training set. The integers (e.g., 92) in the head denote the number of labeled images. Except ST++, the training resolution of other works is larger than us: 512 vs. 321.

credited to a doubled unlabeled batch size. We conjecture that regularizing two strong views with a shared weak view can be regarded as enforcing consistency between these two strong views as well. Intuitively, suppose k_w is the classifier weight of the class predicted by x^w , and (q_{s_1}, q_{s_2}) are features of images (x^{s_1}, x^{s_2}) , then in our adopted cross entropy loss, we maximize $q_j \cdot k_w$ against $\sum_{i=0}^C q_j \cdot k_i$, where $j \in \{s_1, s_2\}$, and k_i is classifier weight of class i . It thus can be deemed that we are also maximizing the similarity between q_{s_1} and q_{s_2} . So the InfoNCE loss [59] is satisfied:

$$\mathcal{L}_{s_1 \leftrightarrow s_2} = -\log \frac{\exp(q_{s_1} \cdot q_{s_2})}{\sum_{i=0}^C \exp(q_j \cdot k_i)}, \quad s.t., j \in \{s_1, s_2\}, \quad (7)$$

where q_{s_1} and q_{s_2} are positive pairs, while all other classifier weights except k_w are negative samples.

Hence, it shares the spirits of contrastive learning [11, 12, 25], which is able to learn discriminative representations and has been proved to be highly meaningful to our task [40, 61].

3.4. Our Holistic Framework: UniMatch

To sum up, we present two key techniques to leverage unlabeled images, namely UniPerb and DusPerb. Our holistic framework (dubbed as UniMatch) that integrates both approaches is illustrated in Figure 2b. The corresponding pseudocode is provided in Algorithm 1. In comparison with FixMatch, two auxiliary feedforward streams are maintained, one for perturbation on features of x^w , and the other for multi-view learning of (x^{s_1}, x^{s_2}) . The final unsupervised term is computed as:

$$\mathcal{L}_u = \frac{1}{B_u} \sum \mathbb{1}(\max(p^w) \geq \tau) \cdot (\lambda \mathbb{H}(p^w, p^{fp}) + \frac{\mu}{2} (\mathbb{H}(p^w, p^{s_1}) + \mathbb{H}(p^w, p^{s_2}))). \quad (8)$$

It is clarified that feature-level and image-level perturbation streams have their own properties and advantages, thus their

Pascal	ResNet-50			ResNet-101		
	1/16	1/8	1/4	1/16	1/8	1/4
SupBaseline	61.2	67.3	70.8	65.6	70.4	72.8
CAC [35]	70.1	72.4	74.0	72.4	74.6	76.3
ST++ [68]	72.6	74.4	75.4	74.5	76.3	76.6
UniMatch 321	74.5	75.8	76.1	76.5	77.0	77.2
SupBaseline	62.4	68.2	72.3	67.5	71.1	74.2
CPS [13]	72.0	73.7	74.9	74.5	76.4	77.7
U ² PL [61]	72.0	75.1	76.2	74.4	77.6	78.7
PS-MT [41]	72.8	75.7	76.4	75.5	78.2	78.7
UniMatch 513	75.8	76.9	76.8	78.1	78.4	79.2
SupBaseline [†]	67.7	71.9	74.5	70.6	75.0	76.5
U ² PL [†] [61]	74.7	77.4	77.5	77.2	79.0	79.3
UniMatch[†] 513	78.1	79.0	79.1	80.9	81.9	80.4

Table 3. Comparison with SOTAs on the **Pascal**. Labeled images are sampled from the *blended* training set. The **321** and **513** denote the training resolution. The fractions in the head denote the proportion of labeled images. We reproduce the RN-50 results of U²PL. †: Prioritizing selecting labeled images from the high-quality set.

loss weights λ and μ are equally set as 0.5. The H in \mathcal{L}_u is a regular cross-entropy loss. The confidence threshold τ is set as 0.95 for all datasets except Cityscapes, where τ is 0.

4. Experiments

4.1. Implementation Details

For a fair comparison with prior works, we mainly adopt DeepLabv3+ [10] based on ResNet [26] as our segmentation model. The ResNet uses an output stride of 16 across all experiments to save memory and speed up training. During training, each mini-batch is composed of 8 labeled images and 8 unlabeled images. The initial learning rate is set as 0.001, 0.005, and 0.004 for Pascal, Cityscapes, and COCO respectively, with a SGD optimizer. The model is trained for 80, 240, and 30 epochs under a poly learning rate scheduler. We assemble the color transformations from ST++ [68] and CutMix [71] to form our \mathcal{A}^s . A raw image is resized between 0.5 and 2.0, cropped, and flipped to obtain its weakly augmented version x^w . The training resolution is set as 321, 801, and 513 for these three datasets. By default, we adopt a channel dropout of 50% probability (`nn.Dropout2d(0.5)` in PyTorch) as our feature perturbation, which is inserted at the intersection of the encoder and decoder.

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

Pascal VOC 2012. The Pascal dataset [17] is originally constructed of 1464 high-quality training images. Later, it is expanded by extra coarse annotations from the SBD [24], resulting in 10582 training images. There are three protocols to select labeled images: (1) **(the most convincing**

Cityscapes	1/16 (186)	1/8 (372)	1/4 (744)	1/2 (1488)
SupBaseline	63.3	70.2	73.1	76.6
PS-MT [41] <small>[CVPR'22]</small>	-	75.8	76.9	77.6
U ² PL [‡] [61] <small>[CVPR'22]</small>	70.6	73.0	76.3	77.2
UniMatch RN-50	75.0	76.8	77.5	78.6
SupBaseline	66.3	72.8	75.0	78.0
CPS [13] <small>[CVPR'21]</small>	69.8	74.3	74.6	76.8
AEL [28] <small>[NeurIPS'21]</small>	74.5	75.6	77.5	79.0
PS-MT [41] <small>[CVPR'22]</small>	-	76.9	77.6	79.1
U ² PL [61] <small>[CVPR'22]</small>	74.9	76.5	78.5	79.1
PCR [65] <small>[NeurIPS'22]</small>	73.4	76.3	78.4	79.1
UniMatch RN-101	76.6	77.9	79.2	79.5

Table 4. Comparison with SOTAs on the **Cityscapes**. ‡: U²PL ResNet-50 results are reproduced on the same splits as ours.

COCO	1/512 (232)	1/256 (463)	1/128 (925)	1/64 (1849)	1/32 (3697)
SupBaseline	22.9	28.0	33.6	37.8	42.2
PseudoSeg [81]	29.8	37.1	39.1	41.8	43.6
PC ² Seg [78]	29.9	37.5	40.1	43.7	46.1
UniMatch	31.9	38.9	44.4	48.2	49.8

Table 5. Comparison with SOTAs on the **COCO** with Xception-65.

one) select from the original high-quality training images, (2) select from the blended 10582 training images, and (3) prioritize the high-quality set, and if not enough, select from the expanded set. Under the first protocol (Table 2), our UniMatch outperforms existing methods tremendously. We even adopt a smaller training size of 321 than most recent works of 512. In addition, for the other two protocols (Table 3), we train UniMatch at two resolutions of 321 and 513. It still gains remarkable improvements over prior works.

Cityscapes. This dataset [15] focuses on urban scenes, consisting of 2975 high-resolution training images and 500 validation images. We follow two techniques in previous SOTA works [13, 28, 41, 61, 65], *i.e.*, sliding window evaluation and online hard example mining (OHEM) loss. Results under ResNet-50 and ResNet-101 are reported in Table 4. Our results are consistently much better than existing best ones.

MS COCO. The COCO dataset [39], composed of 118k/5k training/validation images, is a quite challenging benchmark, containing 81 classes to predict. Therefore, it was rarely explored before in semi-supervised works of segmentation. However, in view of the seemingly saturate performance on the Pascal and Cityscapes, we believe it will be more practical to evaluate our algorithms on this dataset. We adopt exactly the same setting and backbone (Xception-65 [14]) as PseudoSeg [81]. As evidenced by Table 5, our UniMatch significantly surpasses all available methods by 1.4%-4.5%.

Method	92	183	366	732	1464
SOTA Before*	<u>65.2</u>	71.0	74.6	77.3	79.1
FixMatch (Fig 2a)	63.9	<u>73.0</u>	<u>75.5</u>	<u>77.8</u>	<u>79.2</u>
UniPerb (Fig 3a)	72.0	75.8	77.5	79.3	80.1
DusPerb (Fig 3b)	72.1	75.9	78.3	78.1	79.6
UniMatch (Fig 2b)	75.2	77.2	78.8	79.9	81.2
Gain (Δ)	$\uparrow 11.3$	$\uparrow 4.2$	$\uparrow 3.3$	$\uparrow 2.1$	$\uparrow 2.0$

Table 6. Effectiveness of each component on the **Pascal**. *: Borrowed from results in Table 2, that use the same size (321) as us. The better results between our reproduced FixMatch and previous SOTAs are underlined. The Δ is measured against FixMatch.

Method	1/16	1/8	1/4	1/2
FixMatch (Baseline)	72.6	75.7	76.8	78.2
UniMatch (Ours)	75.0	76.8	77.5	78.6
Gain (Δ)	$\uparrow 2.4$	$\uparrow 1.1$	$\uparrow 0.7$	$\uparrow 0.4$

Table 7. Comparison between our UniMatch and our reproduced strong FixMatch baseline on the **Cityscapes** dataset.

Method	1/512	1/256	1/128	1/64	1/32
FixMatch	26.8	32.8	37.7	44.1	47.5
UniMatch	31.9	38.9	44.4	48.2	49.8
Gain (Δ)	$\uparrow 5.1$	$\uparrow 6.1$	$\uparrow 6.7$	$\uparrow 4.1$	$\uparrow 2.3$

Table 8. Comparison between our UniMatch and our reproduced strong FixMatch baseline on the **COCO** dataset.

4.3. Ablation Studies

Unless otherwise specified, we mainly conduct ablation studies on the Pascal dataset extensively with ResNet-101.

Improvement over the FixMatch baseline. We conduct this most important ablation in Table 6, 7, and 8 for all the three benchmarks respectively. It is clear that our UniMatch consistently improves the strong baseline by a large margin.

Individual effectiveness of UniPerb and DusPerb. In Table 6, we first demonstrate that our reproduced FixMatch is a strong competitor against previous SOTA methods. Then built upon FixMatch, both UniPerb and DusPerb facilitate this baseline by a large margin. Lastly, our overall UniMatch that integrates both components achieves the best results.

The improvement of diverse perturbations is non-trivial. Our UniMatch utilizes three views, *i.e.*, one feature perturbation view and dual image perturbation views. We wish to validate that constructing diverse perturbations is beneficial, much better than blindly maintaining three parallel image perturbations. So we design a simple counterpart that adopts three image-level strong perturbation views. As displayed in Table 9, our UniMatch is consistently superior to it, indicating the improvement brought by UniMatch is not credited to blindly increasing views, but the diversity counts.

The improvement of dual-stream perturbation is non-trivial. It might have been noticed that in our DusPerb,




Figure 4. Ablation study on the efficacy of various feature perturbation strategies in our UniPerb method.

Method	92	183	366	732	1464
Dual Image Views	72.1	75.9	78.3	78.1	79.6
Triple Image Views	71.6	76.4	78.4	78.8	79.6
UniMatch	75.2	77.2	78.8	79.9	81.2

Table 9. Ablation study on the non-trivial improvement of diverse perturbations. Our UniMatch is consistently superior to its counterpart which directly uses triple strongly perturbed images as inputs.

Method	92	183	366	732	1464
2× Batch Size	62.5	74.5	77.1	77.8	79.3
2× Epochs	61.8	73.6	76.2	77.6	79.4
DusPerb	72.1	75.9	78.3	78.1	79.6

Table 10. Ablation study on the necessity of dual-stream perturbations, compared with doubling the batch size or training epochs.

the number of unlabeled images in each mini-batch is doubled, since each unlabeled image is strongly perturbed twice. Hence, it might be argued that the improvement is due to larger batch size. Considering this concern, we further carry out an ablation study on the FixMatch with a twice larger batch size (keep the same training iterations) or a twice longer training epochs. As presented in Table 10, although increasing the unlabeled batch size or lengthening training epochs improves the FixMatch baseline in most cases, they are still evidently lagging behind our DusPerb.

The necessity of separating image- and feature-level perturbations into independent streams. PS-MT [41] mixes three levels of perturbations into a single feedforward stream, however, we claim that separating perturbations with different properties into independent streams will enable the model to achieve the targeted invariance more directly, and also avoid a single stream being excessively hard. To confirm it, we inject the dropout on the features of *strongly* perturbed images, forming a stream of hybrid view. As shown is Table 11, one hybrid view is inferior to one image perturbation view. Moreover, we attempt to adopt two hybrid views, but it is still worse than our separate practice in UniPerb.




Figure 5. Ablation study on different values of confidence threshold in our UniMatch method on the **Pascal** dataset. Numbers in the legend denote the number of labeled images.

Method	92	183	366	732	1464
FixMatch	63.9	73.0	75.5	77.8	79.2
Single Hybrid View	63.4	72.8	75.0	76.9	78.7
Dual Hybrid Views	71.8	74.6	76.8	78.4	79.6
UniPerb	72.0	75.8	77.5	79.3	80.1

Table 11. Ablation study on separating image- and feature-level perturbations into independent streams.

IS	FS	92	183	366	732	1464	1/16	1/8	1/4
2	1	75.2	77.2	78.8	79.9	81.2	76.5	77.0	77.2
2	2	75.2	77.7	78.9	79.7	80.7	76.9	77.3	77.9
3	3	75.5	77.0	78.6	79.5	80.5	76.7	77.7	77.3
4	4	75.0	76.6	79.4	79.8	80.6	76.6	77.1	77.5

Table 12. The performance change with respect to the number of image- and feature-level perturbation streams. **IS** stands for image-level stream, while **FS** represents feature-level stream. The first row (IS:2, FS:1) is our UniMatch approach.

More perturbation streams. We also attempt to increase the number of image- and feature-level perturbation streams in Table 12. It is observed that, increasing the perturbation streams does not necessarily result in higher performance. This also indicates that, the two image streams and one feature stream in our UniMatch are well-performed enough. **Other feature perturbation strategies.** We adopt a simplest form of feature perturbation in our method, which is a channel dropout. There are some other options available, such as uniform noise and virtual adversarial training (VAT) [48]. We follow [49] to set the hyper-parameters in these strategies. And all these options are compared in Figure 4. It can be concluded that a channel dropout performs best.

Value of the confidence threshold τ . We ablate this hyper-parameter on the Pascal in Figure 5. It is observed that τ of 0.95 works best for the Pascal.

Locations to insert feature perturbations. Our feature perturbations are injected at the intersection of the encoder and decoder. Previous work [49] also performs perturbations to the input of final classifier. We compare the two locations

Method	WHU-CD				LEVIR-CD			
	5%	10%	20%	40%	5%	10%	20%	40%
S4GAN [47]	18.3 / 96.69	62.6 / 98.15	70.8 / 98.60	76.4 / 98.96	64.0 / 97.89	67.0 / 98.11	73.4 / 98.51	75.4 / 98.62
SemiCDNet [50]	51.7 / 97.71	62.0 / 98.16	66.7 / 98.28	75.9 / 98.93	67.6 / 98.17	71.5 / 98.42	74.3 / 98.58	75.5 / 98.63
SemiCD [3]	65.8 / 98.37	68.1 / 98.47	74.8 / 98.84	77.2 / 98.96	72.5 / 98.47	75.5 / 98.63	76.2 / 98.68	77.2 / 98.72
SupBaseline	48.3 / 97.41	60.7 / 98.03	69.7 / 98.55	69.5 / 98.47	67.5 / 98.12	73.4 / 98.50	75.2 / 98.63	77.7 / 98.79
UniMatch PSPNet	77.5 / 99.06	78.9 / 99.10	82.9 / 99.26	84.4 / 99.32	75.6 / 98.62	79.0 / 98.83	79.0 / 98.84	78.2 / 98.79
SupBaseline	54.1 / 97.56	60.9 / 97.86	68.4 / 98.34	76.2 / 98.87	69.3 / 98.28	76.0 / 98.69	77.6 / 98.79	80.5 / 98.94
UniMatch DeepLab	80.2 / 99.15	81.7 / 99.22	81.7 / 99.18	85.1 / 99.35	80.7 / 98.95	82.0 / 99.02	81.7 / 99.02	82.1 / 99.03

Table 13. Results on the **WHU-CD** [31] and **LEVIR-CD** [9] datasets. Numbers in each cell denote *changed-class IOU* and *overall accuracy*, respectively. The fraction (*e.g.*, 5%) denotes the proportion of labeled images. We try both PSPNet and DeepLabv3+ with ResNet-50.




Figure 6. A typical framework in change detection task. Features extracted by the encoder are subtracted to be fed into the decoder.

Location	92	183	366	732	1464
Decoder-Classifer	64.5	72.2	76.9	77.9	79.5
En-Decoder (Default)	72.0	75.8	77.5	79.3	80.1

Table 14. Ablation study on the location to insert feature perturbations in our UniPerb method.

in Table 14. It is observed that our practice is much better.

4.4. Application to More Segmentation Scenarios

We have validated our UniMatch in common benchmarks of natural images. Here, we further carry out extra experiments in two highly practical and critical scenarios, *i.e.*, remote sensing interpretation and medical image analysis. In both scenarios, unlabeled data is easy and cheap to acquire, while manual annotations are extremely expensive.

Remote Sensing Interpretation. We focus on the change detection task in this scenario, due to its wide application demand and strict labeling requirement. Given a pair of bi-temporal images, *i.e.*, two images for the same region but of different times, the changed regions are required to be highlighted. It can be simply deemed as a binary segmentation problem. A typical framework is illustrated in Figure 6. Following the latest work SemiCD [3], we validate our UniMatch on two popular benchmarks, *i.e.*, WHU-CD [31] and LEVIR-CD [9]. We attempt on two networks, *i.e.* PSPNet and DeepLabv3+, both based on ResNet-50. As shown in Table 13, UniMatch outperforms SemiCD [3] impressively. **Medical Image Analysis.** We follow a recent work [44] to investigate semi-supervised medical image segmentation on the ACDC dataset [4]. As shown in Table 15, our UniMatch

Method	1 case	3 cases	7 cases
SupBaseline	28.5	41.5	62.5
UA-MT [69]	N/A	61.0	81.5
CPS [13]	N/A	60.3	83.3
CNN & Trans [44]	N/A	65.6	86.4
UniMatch (Ours)	85.4	88.9	89.9

Table 15. Comparison with SOTAs on **ACDC** [4] with 1/3/7 labeled cases. There are 70 training cases in total. Results are measured by Dice Similarity Coefficient (DSC) metric averaged on 3 classes.

improves the SOTAs significantly, *e.g.*, +23.3% given 3-labeled cases. Our result of mere 1 labeled case even surpasses others with 3 cases, and is on par with others using 7 cases.

For implementation details of these two scenarios, please refer to our open-sourced code.

5. Conclusion

We investigate the promising role of FixMatch in semi-supervised semantic segmentation. We first present that equipped with proper image-level strong perturbations, a vanilla FixMatch can indeed outperform the SOTAs. Inspired by this, we further strengthen its perturbation practice from two perspectives. On one hand, we unify image- and feature-level perturbations to form a more diverse perturbation space. On the other, we design a dual-stream perturbation technique to fully exploit image-level perturbations. Both components facilitate our baseline significantly. The final method UniMatch improves previous results remarkably in all the natural, medical, and remote sensing scenarios.

Acknowledgements. This work is supported by the NSFC Program (62222604, 62206052, 62192783), CAAI-Huawei MindSpore (CAIXSJLJJ-2021-042A), China Postdoctoral Science Foundation Project (2021M690609), Jiangsu Natural Science Foundation Project (BK20210224), and CCF-Lenovo Bule Ocean Research Fund.

We are sincerely grateful to Boyao Shi (NJU) for polishing the figures, and Zhen Zhao (USYD) for comments.

We also thank the support from SenseTime Research.

References

- [1] Inigo Alonso, Alberto Sabater, David Ferstl, Luis Montesano, and Ana C Murillo. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In *ICCV*, 2021. 3
- [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 *IJCNN*, 2020. 1
- [3] Wele Gedara Chaminda Bandara and Vishal M Patel. Revisiting consistency regularization for semi-supervised change detection in remote sensing images. *arXiv:2204.08454*, 2022. 8
- [4] Olivier Bernard, Alain Lalonde, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, et al. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *TMI*, 2018. 8
- [5] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Mixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. In *ICLR*, 2020. 2, 4
- [6] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019. 2
- [7] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *COLT*, 1998. 3
- [8] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 4
- [9] Hao Chen and Zhenwei Shi. A spatial-temporal attention-based method and a new dataset for remote sensing image change detection. *RS*, 2020. 8
- [10] 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. 5
- [11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2, 5
- [12] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, 2021. 5
- [13] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *CVPR*, 2021. 1, 3, 5, 6, 8
- [14] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *CVPR*, 2017. 6
- [15] 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. 6
- [16] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv:1708.04552*, 2017. 3
- [17] 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
- [18] 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. *PR*, 2022. 3
- [19] Geoff French, Timo Aila, Samuli Laine, Michal Mackiewicz, and Graham Finlayson. Semi-supervised semantic segmentation needs strong, high-dimensional perturbations. In *BMVC*, 2020. 1, 3
- [20] Chengyue Gong, Dilin Wang, and Qiang Liu. Alphamatch: Improving consistency for semi-supervised learning with alpha-divergence. In *CVPR*, 2021. 2
- [21] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *NeurIPS*, 2014. 1, 3
- [22] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *NeurIPS*, 2005. 2
- [23] Dayan Guan, Jiaxing Huang, Aoran Xiao, and Shijian Lu. Unbiased subclass regularization for semi-supervised semantic segmentation. In *CVPR*, 2022. 1, 3
- [24] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, 2011. 5
- [25] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 2, 5
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [27] Ruifei He, Jihan Yang, and Xiaojuan Qi. Re-distributing biased pseudo labels for semi-supervised semantic segmentation: A baseline investigation. In *ICCV*, 2021. 1, 3
- [28] Hanzhe Hu, Fangyun Wei, Han Hu, Qiwei Ye, Jinshi Cui, and Liwei Wang. Semi-supervised semantic segmentation via adaptive equalization learning. In *NeurIPS*, 2021. 1, 3, 6
- [29] Xinyue Huo, Lingxi Xie, Jianzhong He, Zijie Yang, and Qi Tian. Atso: Asynchronous teacher-student optimization for semi-supervised image segmentation. In *CVPR*, 2021. 1
- [30] Jisoo Jeong, Seungeui Lee, Jeesoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. In *NeurIPS*, 2019. 2
- [31] Shunping Ji, Shiqing Wei, and Meng Lu. Fully convolutional networks for multisource building extraction from an open aerial and satellite imagery data set. *TGRS*, 2018. 8
- [32] Zhanghan Ke, Di Qiu, Kaican Li, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. In *ECCV*, 2020. 3
- [33] Chia-Wen Kuo, Chih-Yao Ma, Jia-Bin Huang, and Zsolt Kira. Featmatch: Feature-based augmentation for semi-supervised learning. In *ECCV*, 2020. 2
- [34] Donghyeon Kwon and Suha Kwak. Semi-supervised semantic segmentation with error localization network. In *CVPR*, 2022. 3

- [35] Xin Lai, Zhuotao Tian, Li Jiang, Shu Liu, Hengshuang Zhao, Liwei Wang, and Jiaya Jia. Semi-supervised semantic segmentation with directional context-aware consistency. In *CVPR*, 2021. 3, 5
- [36] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 2
- [37] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop*, 2013. 2
- [38] Junnan Li, Caiming Xiong, and Steven CH Hoi. Comatch: Semi-supervised learning with contrastive graph regularization. In *ICCV*, 2021. 2
- [39] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 6
- [40] Shikun Liu, Shuaifeng Zhi, Edward Johns, and Andrew J Davison. Bootstrapping semantic segmentation with regional contrast. In *ICLR*, 2022. 2, 3, 5
- [41] 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. 2, 3, 4, 5, 6, 7
- [42] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. In *ICLR*, 2021. 1, 3
- [43] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1
- [44] Xiangde Luo, Minhao Hu, Tao Song, Guotai Wang, and Shaoting Zhang. Semi-supervised medical image segmentation via cross teaching between cnn and transformer. In *MIDL*, 2022. 8
- [45] Luke Melas-Kyriazi and Arjun K Manrai. Pixmatch: Unsupervised domain adaptation via pixelwise consistency training. In *CVPR*, 2021. 1, 3
- [46] 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. 3
- [47] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high-and low-level consistency. *TPAMI*, 2019. 1, 3, 8
- [48] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *TPAMI*, 2018. 2, 7
- [49] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *CVPR*, 2020. 1, 3, 7
- [50] Daifeng Peng, Lorenzo Bruzzone, Yongjun Zhang, Haiyan Guan, Haiyong Ding, and Xu Huang. Semicdnet: A semi-supervised convolutional neural network for change detection in high resolution remote-sensing images. *TGRS*, 2020. 8
- [51] Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V Le. Meta pseudo labels. In *CVPR*, 2021. 2
- [52] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan Yuille. Deep co-training for semi-supervised image recognition. In *ECCV*, 2018. 3
- [53] Chuck Rosenberg, Martial Hebert, and Henry Schneiderman. Semi-supervised self-training of object detection models. In *WACV/MOTION*, 2005. 2
- [54] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *NeurIPS*, 2016. 2
- [55] 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. 1, 2, 3
- [56] Nasim Souly, Concetto Spampinato, and Mubarak Shah. Semi supervised semantic segmentation using generative adversarial network. In *ICCV*, 2017. 1, 3
- [57] Yihe Tang, Weifeng Chen, Yijun Luo, and Yuting Zhang. Humble teachers teach better students for semi-supervised object detection. In *CVPR*, 2021. 1, 3
- [58] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 2
- [59] Aaron Van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018. 5
- [60] Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, Zhen Wu, and Jindong Wang. Freematch: Self-adaptive thresholding for semi-supervised learning. *arXiv:2205.07246*, 2022. 2
- [61] 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, 5, 6
- [62] Junfei Xiao, Longlong Jing, Lin Zhang, Ju He, Qi She, Zongwei Zhou, Alan Yuille, and Yingwei Li. Learning from temporal gradient for semi-supervised action recognition. In *CVPR*, 2022. 1, 3
- [63] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. In *NeurIPS*, 2020. 2
- [64] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *CVPR*, 2020. 2
- [65] Hai-Ming Xu, Lingqiao Liu, Qiuchen Bian, and Zhen Yang. Semi-supervised semantic segmentation with prototype-based consistency regularization. In *NeurIPS*, 2022. 5, 6
- [66] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *ICCV*, 2021. 1, 3
- [67] Yinghao Xu, Fangyun Wei, Xiao Sun, Ceyuan Yang, Yujun Shen, Bo Dai, Bolei Zhou, and Stephen Lin. Cross-model pseudo-labeling for semi-supervised action recognition. In *CVPR*, 2022. 1, 3
- [68] 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, 3, 5

- [69] Lequan Yu, Shujun Wang, Xiaomeng Li, Chi-Wing Fu, and Pheng-Ann Heng. Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation. In *MICCAI*, 2019. 8
- [70] 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, 3
- [71] 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. 2, 3, 5
- [72] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. In *NeurIPS*, 2021. 2
- [73] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *CVPR*, 2018. 1
- [74] Jianrong Zhang, Tianyi Wu, Chuanghao Ding, Hongwei Zhao, and Guodong Guo. Region-level contrastive and consistency learning for semi-supervised semantic segmentation. In *IJCAI*, 2022. 3
- [75] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, and Fang Wen. Robust mutual learning for semi-supervised semantic segmentation. *arXiv:2106.00609*, 2021. 3
- [76] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *CVPR*, 2018. 3
- [77] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. 1
- [78] Yuanyi Zhong, Bodi Yuan, Hong Wu, Zhiqiang Yuan, Jian Peng, and Yu-Xiong Wang. Pixel contrastive-consistent semi-supervised semantic segmentation. In *ICCV*, 2021. 3, 5, 6
- [79] 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. 3
- [80] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin D Cubuk, and Quoc V Le. Rethinking pre-training and self-training. In *NeurIPS*, 2020. 2
- [81] Yulian 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. 1, 3, 6