

# Beyond Pixels: Semi-Supervised Semantic Segmentation with a Multi-scale Patch-based Multi-Label Classifier

Prantik Howlader<sup>1</sup>, Srijan Das<sup>2</sup>, Hieu Le<sup>3</sup>, and Dimitris Samaras<sup>1</sup>

<sup>1</sup> Stony Brook University, New York, USA

<sup>2</sup> University of North Carolina at Charlotte, NC, USA

<sup>3</sup> EPFL, Lausanne, Switzerland


**Abstract.** Incorporating pixel contextual information is critical for accurate segmentation. In this paper, we show that an effective way to incorporate contextual information is through a patch-based classifier. This patch classifier is trained to identify classes present within an image region, which facilitates the elimination of distractors and enhances the classification of small object segments. Specifically, we introduce **Multi-scale Patch-based Multi-label Classifier** (MPMC), a novel plug-in module designed for existing semi-supervised segmentation (SSS) frameworks. MPMC offers patch-level supervision, enabling the discrimination of pixel regions of different classes within a patch. Furthermore, MPMC learns an adaptive pseudo-label weight, using patch-level classification to alleviate the impact of the teacher’s noisy pseudo-label supervision on the student. This lightweight module can be integrated into any SSS framework, significantly enhancing their performance. We demonstrate the efficacy of our proposed MPMC by integrating it into four SSS methodologies and improving them across two natural image and one medical segmentation dataset, notably improving the segmentation results of the baselines across all the three datasets. Code is available at: <https://github.com/cvlab-stonybrook/Beyond-Pixels>

**Keywords:** Segmentation · Semi-supervised · Patches

## 1 Introduction

Supervised segmentation requires a large amount of pixel level annotations, which is laborious and time-consuming [53–56]. This paper focuses on semi-supervised segmentation (SSS) [1, 20], which alleviates the reliance on labeled annotations by leveraging a large quantity of unlabeled images, accompanied by a limited number of manually labeled images.

The fundamental task in segmentation is a per-pixel classification. Thus contextual information of a pixel is important to improve its performance [31, 32, 68]. Context becomes more critical in SSS settings, which uses limited labeled data making the models more prone to overfitting. We note that recent approaches in SSS methods which predominantly leverage consistency regularization with



**Fig. 1: Segmentations and Multi-label classifications of image patches (Cityscapes)** After 20 epochs, (b) and (c) are the segmentations and classes in the segmentations respectively, (d) is classes predicted by multi-label classifier in MPMC. (f) and (g) are segmentations by Baseline (AugSeg [65]) and Baseline + MPMC at the end of training. White boxes represent regions where MPMC improves the baseline, while orange boxes represent regions where segmentation network predicts classes not present in the patch.

pseudo-labeling techniques [7, 14, 20, 22, 36, 38, 50, 57, 69] do not exploit contextual information explicitly.

We hypothesize that *patches* are the simplest way to determine the context around pixels. The key observation is that the contextual information of an image patch are the classes present in it. For example, in the image patch in the first row of Fig. 1 (a) the classes present are car, road, vegetation and terrain. Thus, training a network to not only classify each pixel in an image but also predict the classes in the patch surrounding it, is a straight forward way to introduce contextual information in SSS pipeline. Moreover, utilizing multi-scale patch-level supervision can effectively encapsulate context at varying scales and distinguish pixels across different classes. To this end, we introduce a novel plug-in module, termed as **Multi-scale Patch-based Multilabel Classifier** (MPMC). This module achieves enhanced discrimination of pixel regions between classes through additional multi-scale patch-level supervision. Notably, we observe that when training our baseline SSS model together with MPMC in limited labeled data scenario, the patch classifier exhibits greater resilience to noise than the segmentation network as illustrated in Fig. 1 (b) and (c). A key observation is when MPMC and baseline SSS model are trained together, the patch classifier provides patch-level labels, identifying the classes present within a region. This facilitates the elimination of distractors (first row in Fig. 1 (b), (c) and (d), while the segmentation network classifies pixels as truck which is not present in the patch, the multi-label classifier does not predict it) and enhances the classification of small object segments by the baseline method.

Another important issue plaguing SSS methods is noisy pseudo-labels generated by the teacher for unlabeled images to supervise the student. This can lead to the student model overfitting to these noisy labels, consequently impairing its performance. Recent approaches have attempted to mitigate the impact of noise in pseudo-labels through various strategies, such as implementing a hard threshold [14, 23, 38], applying dynamic re-weighting to enhance the influence of reliable pixels [20], and employing negative data mining by archiving unreliable pixels in a memory bank [50]. In contrast to these approaches, we use the MPMC predictions to learn an adaptive weight that reduces the weight of noisy pseudo-labels in the teacher-student framework. Our approach hinges on the underlying observation that *MPMC exhibits greater uncertainty (lower confidence) for noisy labels in a patch* (see supplementary for details) .

In practice, MPMC is a parameterized module that can be plugged in the initial layers of the visual model, where it leverages the low-level pixel information to classify the classes present in each patch. This patch-level classification strategy forces the visual model to learn and preserve discriminative representations of classes, which are often prone to dilution in the deeper layers of the model. We demonstrate the efficacy of our proposed parameterized module, MPMC by integrating it into four SSS methodologies—UniMatch [57], AugSeg [65], AEL [20], and U2PL [50] - notably improving performance across all data partitions in two natural (Cityscapes and Pascal VOC) and one medical (ACDC) dataset. To summarize, our contributions are:

1. We introduce a novel plug-in module, called MPMC that incorporates an auxiliary multi-scale multi-label classification task, enhancing pixel region discrimination across different classes via additional patch-level supervision.
2. MPMC consists of two adaptive learning weights that reduces the influence of noisy pseudo-labels from the teacher model.
3. We demonstrate MPMC’s robustness by integrating in four SSS techniques, achieving consistent improvements across all protocols on Pascal VOC and Cityscapes. Notably, our method also shows improvement in medical image dataset ACDC, which indicates that our method is applicable beyond natural images.

## 2 Related Work

Semi-supervised learning (SSL) has been extensively researched in recent years. Popular SSL approaches include consistency regularization [2, 10, 11, 14, 38, 43, 58], entropy minimization [5, 16], and pseudo-labeling [24, 28, 46]. In this work, our emphasis is placed on the domains of pseudo-labeling and consistency regularization.

**Pseudo labeling:** Pseudo-labeling methods [28, 45] alongside self-training techniques [35, 59] are designed to leverage labeled image data to produce pseudo-labels for unlabeled data. An inherent problem that plagues pseudo-labeling is noise, thus most approaches [42, 44, 46, 69] devise strategies to refine the pseudo-labels. PseudoSeg [69], for instance, aims at enhancing pseudo-label accuracy

using attention mechanisms based on grad-cam [44]. In pseudo-label selection, FixMatch [46] adopts a confidence threshold, whereas UPS [42] relies on uncertainty measures. ST++ [58] focuses on selecting unlabeled images more likely to yield accurate pseudo-pixels. Emerging approaches like U2PL [50] explore the idea of learning reliable pseudo-labels from their less reliable counterparts. CFCG [29] introduced a novel approach that combines cross-fusion and contour guidance for pseudo-label improvement. We differ from these methods by explicitly introducing patch-level contextual information in the SSS methods.

**Consistency Regularization:** This approach enforces the model to have consistent predictions on different views of the unlabeled images. UniMatch [57] uses dual perturbations and feature perturbations to force the network to learn more consistent information. ICT [49] uses Mixup [64] augmentation for consistency regularization. Recent methods [14, 58, 62] use Cutout [9], Cutmix [63], Classmix [37] as strong data augmentation. Consistency Regularization is used with pseudo-labeling techniques in Mixmatch [4] and TC-SSL [67]. In our work, we use a set of strong and weak data augmentations [1, 65] for generating pseudo-labels.

**Multi-label classification:** This approach classifies an image into more than one label present in the image. Most works in this field explore finding the correlation between the labels and the images [21, 39, 61]. Recent works mainly focus on solving data imbalance issues [17, 51]. We use multi-label classification to help in segmentation. Some works [15, 40] use multi-label classification for segmentation in weakly supervised scenarios, for improved CAM-based pseudo-segmentation. However, we are using multi-label classification in a semi-supervised setting to introduce patch-level contextual information and also use the classifier confidence to reduce the impact of noisy pseudo-labels.


### 3 Semi-Supervised Segmentation Pipeline

Given a labeled set  $\mathcal{D}^l = \{(x^l, y^l)\}$  and an unlabeled set  $\mathcal{D}^u = \{x^u\}$ , the goal of semi-supervised semantic segmentation is to develop a model that effectively utilizes both labeled and unlabeled data. This section commences with an overview of the foundational framework for semi-supervised segmentation in Section 3.1 followed by an overview of our proposed MPMC in Section 3.2. The section concludes with an elaboration of MPMC’s functionality (Section 3.3).

#### 3.1 Preliminary

Teacher-Student frameworks leverage weak-to-strong consistency regularization to leverage unlabeled data. As depicted in Fig. 2(a), each unlabeled image  $x^u$  is simultaneously perturbed by two operators, i.e., weak perturbation  $\mathcal{A}^w$  (cropping, random flip, etc), and strong perturbation  $\mathcal{A}^s$  (color jitter, gaussian blur, contrast, etc). The framework’s objective function is a convex combination of supervised loss  $\mathcal{L}_l$  for the labeled data and unsupervised loss  $\mathcal{L}_u$  for the unlabeled data as:

$$\mathcal{L} = \mathcal{L}_l + \alpha \mathcal{L}_u \quad \text{where, } \mathcal{L}_u = \frac{1}{|\mathcal{D}^u|} \sum \mathbb{1}(p^w \geq th) \mathcal{E}(p^w, p^s) \quad (1)$$



**Fig. 2: Overall Pipeline of our novel multi-label classification based semantic segmentation:** (a) **Left:** End-to-end Teacher-Student Pipeline with our novel method Multi-scale Patch-based Multi-label Classifier (MPMC). (b) **Right:** For unlabeled images, The teacher MPMC extracts features from a layer in the segmentation model’s encoder to classify the feature’s receptive field patch. The confidence of MPMC for a class in a patch is used to calculate two adaptive weights  $\lambda_s$  and  $\lambda_m$  which is used to reduce the influence of noisy predictions in that patch from the teacher to train both the student segmentation network and MPMC.

where,  $\mathcal{L}_l$  is the standard cross-entropy loss, comparing the model’s predictions with ground truth. While  $\mathcal{L}_u$  aims to regularize the predictions for an unlabeled sample so that they remain consistent under both weak and strong augmentations ( $p^w$  and  $p^s$ ). Here,  $th$  is a predefined confidence threshold to filter noisy pseudo labels and  $\mathcal{E}$  minimizes the entropy between two probability distributions.

### 3.2 Overview

Fig. 2(a) illustrates the comprehensive training process of our proposed MPMC, which is tailored to enhance the discrimination of patches among different classes and to mitigate the impact of noisy predictions transmitted from the teacher to the student. MPMC can be integrated into both teacher and student models within the teacher-student framework. It is composed of a classifier that performs multi-label classifications of the multi-scale patch representations of the input image. The application of MPMC spans two distinct scenarios: 1) For labeled images, MPMC assigns classes to each patch representation based on the classes present within a patch. The multi-scale nature of these patch representations and their classification enhances the segmentor’s ability to discriminate classes by capturing enhanced contextual information related to variations in shape and size. 2) In the case of unlabeled images, besides patch classification, MPMC computes a pseudo-label weight map to alleviate the impact of the teacher’s noisy pseudo-label supervision on the student. This involves assigning weights to the pseudo-labels within each patch in consistency learning framework, which are determined by the confidence of the MPMC’s predictions for each class in the patch, to enhance the reliability and accuracy of the pseudo-labeling process.

### 3.3 Multi-scale Patch-based Multi-label Classifier (MPMC)

Our proposed MPMC module is specifically designed to address the shortcomings of SSS pipeline. MPMC can be plugged in at any layer within the teacher-student network, both within the teacher model, denoted as  $g$ , and within the student model, denoted as  $\hat{g}$ , they share the same parameters. Below, we detail the training of MPMC on labeled images and then on unlabeled images.

**MPMC on Labeled Images.** As illustrated in Fig. 2, MPMC incorporates patch-supervision into the segmentation framework, by training a multi-label classifier that takes as input a single pixel in the feature map of the  $\tau^{th}$  layer of the encoder. This single pixel-representation encodes the information from a patch (receptive field) in the original image. Further, to encapsulate context at varying scales we perform average pooling on this features map with kernels of different sizes before performing multi-label classification. Let us define a patch of image and the corresponding segmentation in a receptive field  $r$  as  $\{x_r^l, y_r^l\} \in \mathcal{D}^l$ . Note that  $y_r^l$  is a multi-label target where the classes present in the receptive field  $r$  are positive classes, and the classes not present are negative classes.

The input to the student MPMC  $\hat{g}$  are the features extracted from the  $\tau^{th}$  layer of the encoder in student segmentation network  $\hat{f}$ , denoted by  $z_r = \hat{f}_\tau(\mathcal{A}_w(x_r^l))$ . MPMC captures the contextual information within a patch at different scales,  $S_i$ ,  $i \in \{0, 1\}$ , where  $s_0$  and  $s_1$  are average pool kernel sizes. For all our experiments  $s_0$  and  $s_1$  are set to  $7 \times 7$  and  $19 \times 19$ . These pooling operations are performed with appropriate padding to maintain uniformity in the shape of the output features. MPMC  $\hat{g}$  is composed of the concatenation of these multi-scale pooled features, followed by a convolutional network  $\mathcal{M}$  and a final linear layer. Thus, the output of the MPMC  $\hat{g}$  for the input patch  $x_r^l$  is given by:

$$q^l = \sigma(W^T \mathcal{M}(\text{concat}(z_r, z_r^{S_0}, z_r^{S_1})) + b) \quad (2)$$

where  $W$  and  $b$  are linear layer weights and bias,  $\sigma$  is the sigmoid activation and  $q^l = [q_1^l, \dots, q_C^l]^T \in \mathbb{R}^C$  are the predicted probabilities for  $C$  classes.

While binary cross-entropy loss is applicable to train MPMC, it indiscriminately weighs all positive and negative classes equally. Considering the receptive field of a feature from the  $\tau^{th}$  layer of the encoder in the segmentation network  $\hat{f}$  covers only a small patch relative to the input image's size, the prevalence of negative classes within a patch significantly exceeds that of positive classes. This imbalance influences the binary cross-entropy loss to favor negative classes [51]. Inspired by long-tailed multi-label classification [30, 51] which mitigate such imbalances by diminishing the weight of negative classes, we adapt an asymmetric loss [41], a focal loss variant that assigns distinct  $\gamma$  values to positive and negative classes. For an image patch  $x_r^l$  and the predicted class probabilities  $q^l$  using our framework, the asymmetric focal loss is expressed as:

$$l_r^M = \frac{1}{C} \sum_{c=1}^C \begin{cases} (1 - q_c^l)^{\gamma^+} \log(q_c^l), & y_{r,c}^l = 1. \\ (q_c^l)^{\gamma^-} \log(1 - q_c^l), & y_{r,c}^l = 0. \end{cases} \quad (3)$$

In our experiments, we set  $\gamma^+ = 0$  and  $\gamma^- = 1$ . Subsequently, MPMC  $\hat{g}$  undergoes training with multi-label supervision  $\mathcal{L}_l^M$ , formulated as:

$$\mathcal{L}_l^M = \frac{1}{|D^l|} \sum_{(x^l, y^l) \in D^l} \sum_{r=1}^R l_r^M \quad (4)$$

where  $R$  is the number of receptive fields at  $\tau^{th}$  layer of the segmentor. This auxiliary task of patch classification augments the network’s proficiency in learning and preserving class features, a critical function particularly pertinent as the fidelity of these features tends to diminish in the network’s deeper layers.

**MPMC on Unlabeled images.** In semi-supervised teacher-student framework, for unlabeled images, the prediction of the teacher network  $f$  for the weak augmented view  $\mathcal{A}^w(x^u)$  is used to supervise the student network  $\hat{f}$  for the strong augmented view  $\mathcal{A}^s(x^u)$ . In the early iterations of semi-supervised learning, the teacher’s predictions are inherently noisy. Common practices [13, 20, 33, 46, 70], involve employing pixel-wise confidence score thresholds ( $th$ ) to filter out unreliable pseudo labels. However, a high confidence threshold does not entirely eliminate noise due to the teacher’s initial miscalibration [42], allowing some noisy predictions to have high confidence. To refine this process, our approach does not solely depend on confidence-thresholded teacher segmentation predictions for reliable pseudo-label extraction. Instead, we leverage the teacher MPMC  $g$ ’s confidence for a patch to adjust the weights of the deemed reliable pseudo-labels within that patch. The underlying principle of our methodology is that *MPMC exhibits greater uncertainty (lower confidence) for noisy labels*.

We employ the teacher MPMC  $g$ ’s predictions to define two adaptive learning weights  $\lambda_m \in \mathbb{R}^{R \times C}$  and  $\lambda_s \in \mathbb{R}^{W \times H}$ , where  $R$  denotes the number of patches, sized according to the receptive field within the image,  $W$  and  $H$  represents the width and height of the input image, respectively. The learning weight  $\lambda_m$  is utilized to assign class weights to each patch for training the student MPMC, whereas  $\lambda_s$  is applied to diminish the impact of noisy predictions during training the student segmentation network.

For an unlabeled image  $x^u$ , consider a patch at the receptive field  $r$ , denoted as  $x_r^u$ . The teacher and student segmentation networks yield predictions  $p^w = f(\mathcal{A}^w(x_r^u))$  and  $p^s = \hat{f}(\mathcal{A}^s(x_r^u))$ , respectively. Simultaneously, the teacher and student MPMC predict  $q_r^u = g(f_\tau(\mathcal{A}^w(x_r^u)))$  and  $\hat{q}_r^u = \hat{g}(\hat{f}_\tau(\mathcal{A}^s(x_r^u)))$  using features from the  $\tau^{th}$  layer of the encoder. Let us define the confidence of teacher MPMC  $g$  across all classes in patch  $r$  as  $\gamma = \text{softmax}(q_r^u)$ , with  $\gamma \in \mathbb{R}^C$ . The weight of the patch  $r$ ,  $\lambda_{m,r}$  is determined by  $\gamma$ . For each pixel in patch  $r$  associated with class  $c$ ,  $\lambda_s$  is assigned the weight  $\gamma_c$ . This approach assigns weights to pseudo-labels and pixels based on the teacher MPMC’s confidence, enhancing the accuracy of the learning process.

Thus, the control parameters  $\lambda_s$  and  $\lambda_m$  are integrated to the unsupervised losses to train the student segmentation network and MPMC. The unsupervised



segmentation loss  $L_u$  mentioned in equation 1 to train the student segmentation network  $\hat{f}$  is re-formulated as:

$$\mathcal{L}_u = \frac{1}{|\mathcal{D}^u|} \sum_{x^u \in \mathcal{D}^u} \mathbb{1}(p^w \geq th) \lambda_s \mathcal{E}(p^w, p^s) \quad (5)$$

where  $\mathcal{E}$  is the standard cross entropy loss and  $th$  is the predefined segmentation confidence threshold. Finally, the unsupervised multi-label loss  $\mathcal{L}_u^M$  to train the student MPMC  $\hat{g}$  is given by:

$$\mathcal{L}_u^M = \frac{1}{|\mathcal{D}^u|} \sum_{x^u \in \mathcal{D}^u} \sum_{r=1}^R \frac{1}{C} \sum_{i=1}^C \lambda_{m,r}^i \mathcal{E}^+(q_r^u, \hat{q}_r^u) \quad (6)$$

where  $\mathcal{E}^+$  is the binary cross entropy loss. Therefore, the overall loss function is defined as:

$$\mathcal{L} = \mathcal{L}_l + \mathcal{L}_l^M + \alpha \mathcal{L}_u + \beta \mathcal{L}_u^M \quad (7)$$

where  $\alpha$  and  $\beta$  controls the contribution of  $\mathcal{L}_u$  and  $\mathcal{L}_u^M$ . In all our experiments  $\alpha$  and  $\beta$  are set to 0.1 and 0.25 respectively.

## 4 Experiments

### 4.1 Setup

**Datasets:** **PASCAL VOC 2012** [12] is a standard semantic segmentation dataset with 20 semantic classes and a background class. The training, validation, and testing sets consist of 1464, 1449, and 1556 images (*Classic*), respectively. Following [7, 65], we also include the additional augmented dataset [18] (*Blender*), which includes a collection of 10582 training images. We adopt the same partition protocols in [7, 65] to evaluate our method in both *Classic* and *Blender* sets. **Cityscapes** [8] dataset is designed for urban scene understanding, consisting of 30 classes, of which only 19 classes are used for scene parsing evaluation. It has 2975, 500, and 1525 images for training, validation, and testing, respectively. **ACDC** [3] dataset contains 200 annotated short-axis cardiac cine-MR images from 100 patients. The segmentation masks consists of classes: left ventricle (LV), myocardium (Myo), and right ventricle (RV). 140 images from 70 patients, 60 images from 30 patients are randomly selected for training and validation respectively.

**Implementation Details:** For fair comparison, we adopt DeepLabv3+ [6] as the decoder in all of our experimental setups, and compare with both ResNet-101 and ResNet-50 [19] as the backbone architecture. In all our experiments with Cityscapes, Pascal, and ACDC datasets, the training resolution is 801, 513, and 224 respectively. Note, for maintaining fair comparison, our method “UniMatch [57] + MPMC” uses the training resolution of 321 which is used by UniMatch [57] when training on the Classic set of Pascal VOC. Further when training a MPMC



integrated baseline, we use the same weak and strong augmentations as used by the corresponding baselines. The convolutional network in MPMC is comprised of three convolution blocks. Each convolution block consists of a stack of two  $3 \times 3$  convolutions followed by a  $1 \times 1$  convolution. For all our experiments MPMC is integrated to the first layer of the ResNet Encoder.

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

We conducted our experiments on PASCAL VOC 2012 and Cityscapes which are two popular natural image benchmarks, and a medical dataset ACDC. MPMC is integrated in four semi supervised methods UniMatch [57], AEL [20], U2PL [50] and AugSeg [65]. Note, UniMatch [57] is a consistency regularization based method. We observe that MPMC consistently improves over their baseline performance across all partition protocols on all datasets.

**Results on PASCAL VOC 2012 Dataset.** Table 1 shows the comparison with other SOTA methods for both the *Classic* and *Blender* sets. MPMC integrated to the four baseline methods improves them consistently across all data partitions for both *Classic* and *Blender* sets. We find that our method brings the most improvement for the least labeled data partition ( $\frac{1}{16}$ ), improving AugSeg, AEL, U2PL and UniMatch by **1.8%**, **2.5%**, **2.0%** and **2.1%** respectively in *Classic* with a ResNet-101 encoder.

**Results on Cityscapes Dataset.** Table 2 shows the comparison with other SOTA methods. MPMC integrated to the four baseline methods improves them consistently across all data partitions. We observe that similar to results in PASCAL VOC 2012 dataset, our method brings the most improvement for the least labeled data partition ( $\frac{1}{16}$ ), improving AugSeg, AEL, U2PL and UniMatch by **1.7%**, **2.1%**, **2.3%** and **1.9%** respectively in *Classic* for ResNet-101 encoder.


**Results on ACDC Dataset.** Table 3 shows the comparison with other SOTA methods. We observe that when MPMC is integrated to the four baseline methods, it significantly improves the performance for the “3 cases” data partition which has the least labeled images ( $< 10\%$  of labeled images). Note we use ResNet-101 for all experiments in this dataset.

The consistent improvement in SSS performance in both natural and medical datasets and its ability to integrate in different SSS methods substantiates the robustness of MPMC. Moreover, the notable improvements achieved in scenarios with limited data shows MPMC’s vital efficacy within SSS pipelines.

**Qualitative Results** In Fig. 3 we visualise the results of integration of our method (MPMC) to AugSeg [65]. Introducing patch-level supervision with the integration of MPMC, leads the SSS model to better segment confusing classes. Among comparisons we show that MPMC leads to better segmentation of pixels regions of confusing classes like bus, train and truck. These classes are often misclassified not only because of sharing visual similarity but also having low training instances.

**Table 1:** Quantitative results of different semi-supervised segmentation methods on Pascal VOC classic and blender set. We report Mean IoU under various partition protocols and show the improvements ( $\Delta$ ) over corresponding baseline.

| Method                             | Classic |        |        |        |        | Blender |        |        |
|------------------------------------|---------|--------|--------|--------|--------|---------|--------|--------|
|                                    | 1/16    | 1/8    | 1/4    | 1/2    | Full   | 1/16    | 1/8    | 1/4    |
| <b>ResNet-50</b>                   |         |        |        |        |        |         |        |        |
| <i>Supervised</i>                  | 44.0    | 52.3   | 61.7   | 66.7   | 72.9   | -       | -      | -      |
| PC <sup>2</sup> Seg [66] [CVPR'21] | 56.9    | 64.6   | 67.6   | 70.9   | 72.3   | -       | -      | -      |
| PseudoSeg [69] [ICLR'21]           | 56.9    | 64.6   | 67.6   | 70.9   | 72.2   | -       | -      | -      |
| ST++ [58] [CVPR'22]                | -       | -      | -      | -      | -      | 72.6    | 74.4   | 75.4   |
| AugSeg [65] + MPMC [CVPR'21]       | 64.2    | 72.1   | 76.1   | 77.4   | 78.8   | 77.2    | 78.2   | 78.2   |
| ( $\Delta$ )                       | (+1.7)  | (+1.6) | (+1.5) | (+1.1) | (+0.9) | (+1.8)  | (+1.6) | (+1.5) |
| AEL [20] [NeurIPS'21]              | 62.9    | 64.1   | 70.3   | 72.7   | 74.0   | 74.1    | 76.1   | 77.9   |
| AEL [20] + MPMC ( $\Delta$ )       | 65.0    | 66.1   | 72.1   | 74.2   | 75.3   | 76.1    | 78.0   | 79.6   |
| ( $\Delta$ )                       | (+2.1)  | (+2.0) | (+1.8) | (+1.5) | (+1.3) | (+2.0)  | (+1.9) | (+1.7) |
| U2PL [50] [CVPR'22]                | 63.3    | 65.5   | 71.6   | 73.8   | 75.1   | 74.7    | 77.4   | 77.5   |
| U2PL [50] + MPMC ( $\Delta$ )      | 65.2    | 67.4   | 73.3   | 75.2   | 76.2   | 76.5    | 79.1   | 79.1   |
| ( $\Delta$ )                       | (+1.9)  | (+1.9) | (+1.7) | (+1.4) | (+1.1) | (+1.8)  | (+1.7) | (+1.6) |
| UniMatch [57] [CVPR'23]            | 71.9    | 72.5   | 76.0   | 77.4   | 78.7   | 78.1    | 79.0   | 79.1   |
| UniMatch [57] + MPMC ( $\Delta$ )  | 73.5    | 74.0   | 77.4   | 78.5   | 79.4   | 79.6    | 80.4   | 80.4   |
| ( $\Delta$ )                       | (+1.6)  | (+1.5) | (+1.4) | (+1.1) | (+0.7) | (+1.5)  | (+1.4) | (+1.3) |
| <b>ResNet-101</b>                  |         |        |        |        |        |         |        |        |
| <i>Supervised</i>                  | 45.1    | 55.3   | 64.8   | 69.7   | 73.5   | 70.6    | 75.0   | 76.5   |
| CPS [7] [CVPR'21]                  | 64.1    | 67.4   | 71.7   | 75.9   | -      | 72.2    | 75.8   | 77.6   |
| PS-MT [33] [CVPR'22]               | 65.8    | 69.6   | 76.6   | 78.4   | 80.0   | 75.5    | 78.2   | 78.7   |
| PCR [52] [NeurIPS'22]              | 70.1    | 74.7   | 77.2   | 78.5   | 80.7   | 78.6    | 80.7   | 80.8   |
| DAW [47] [NeurIPS'23]              | 74.8    | 77.4   | 79.5   | 80.6   | 81.5   | 78.5    | 78.9   | 79.6   |
| CFCG [29] [ICCV'23]                | -       | -      | -      | -      | -      | 77.4    | 79.4   | 80.4   |
| AugSeg [65] [CVPR'21]              | 71.1    | 75.5   | 78.8   | 80.3   | 81.4   | 79.3    | 81.5   | 80.5   |
| AugSeg [65] + MPMC ( $\Delta$ )    | 72.9    | 77.8   | 79.9   | 81.3   | 82.1   | 80.8    | 82.5   | 81.4   |
| ( $\Delta$ )                       | (+1.8)  | (+2.3) | (+1.1) | (+1.0) | (+0.7) | (+1.5)  | (+1.0) | (+0.9) |
| AEL [20] [NeurIPS'21]              | 66.1    | 68.3   | 71.9   | 74.4   | 78.9   | 77.2    | 77.6   | 78.1   |
| AEL [20] + MPMC ( $\Delta$ )       | 68.6    | 70.3   | 73.5   | 76.1   | 80.1   | 79.2    | 79.3   | 79.2   |
| ( $\Delta$ )                       | (+2.5)  | (+2.0) | (+1.6) | (+1.7) | (+1.2) | (+2.0)  | (+1.7) | (+1.1) |
| U2PL [50] [CVPR'22]                | 68.0    | 69.2   | 73.7   | 76.2   | 79.5   | 77.2    | 79.0   | 79.3   |
| U2PL [50] + MPMC ( $\Delta$ )      | 70.0    | 70.8   | 75.3   | 77.8   | 80.5   | 79.2    | 80.7   | 80.3   |
| ( $\Delta$ )                       | (+2.0)  | (+1.6) | (+1.6) | (+1.6) | (+1.0) | (+2.0)  | (+1.7) | (+1.0) |
| UniMatch [57] [CVPR'23]            | 75.2    | 77.2   | 78.8   | 79.9   | 81.2   | 80.9    | 81.9   | 80.4   |
| UniMatch [57] + MPMC ( $\Delta$ )  | 77.3    | 78.6   | 79.8   | 80.8   | 81.7   | 82.5    | 83.1   | 81.5   |
| ( $\Delta$ )                       | (+2.1)  | (+1.4) | (+1.0) | (+0.9) | (+0.5) | (+1.6)  | (+1.2) | (+1.1) |



**Fig. 3: Qualitative Results on Cityscapes dataset:** (a) is the original image, (b) is the Ground Truth, (c) are segmentations generated by AugSeg [65] compared to (d) which are segmentations generated when our method (MPMC) is integrated to AugSeg. White boxes show the areas where our method improves the baseline.

**Table 2: Quantitative results of different semi-supervised segmentation methods on Cityscapes validation set.** We report Mean IoU under various partition protocols and show the improvements ( $\Delta$ ) over the corresponding baseline.

| Method               |              |           | ResNet-50   |             |             |             | ResNet-101  |             |             |             |  |
|----------------------|--------------|-----------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|--|
|                      |              |           | 1/16        | 1/8         | 1/4         | 1/2         | 1/16        | 1/8         | 1/4         | 1/2         |  |
| <i>Supervised</i>    |              |           | 63.34       | 68.73       | 74.14       | 76.62       | 66.3        | 72.8        | 75.0        | 78.0        |  |
| CPS [7]              | [CVPR'21]    | [CVPR'22] | 69.79       | 74.39       | 76.85       | 78.64       | 69.8        | 74.3        | 74.6        | 76.8        |  |
| PS-MT [33]           |              |           | -           | 75.76       | 76.92       | 77.64       | -           | 76.9        | 77.6        | 79.1        |  |
| PCR [52]             |              |           | -           | -           | -           | -           | 73.4        | 76.3        | 78.4        | 79.1        |  |
| CFCG [29]            |              |           | 76.1        | 78.9        | 79.3        | 80.1        | 77.8        | 79.6        | 80.4        | 80.9        |  |
| AugSeg [65]          | [CVPR'21]    |           | 73.7        | 76.4        | 78.7        | 79.3        | 75.2        | 77.8        | 79.6        | 80.4        |  |
| AugSeg [65] + MPMC   |              |           | <b>75.3</b> | <b>77.7</b> | <b>79.8</b> | <b>80.3</b> | <b>76.9</b> | <b>79.2</b> | <b>80.9</b> | <b>81.3</b> |  |
| ( $\Delta$ )         |              |           | (+1.6)      | (+1.3)      | (+1.1)      | (+1.0)      | (+1.7)      | (+1.4)      | (+1.3)      | (+0.9)      |  |
| AEL [20]             | [NeurIPS'21] |           | 68.2        | 72.7        | 74.9        | 77.5        | 74.5        | 75.6        | 77.5        | 79.0        |  |
| AEL [20] + MPMC      |              |           | <b>70.5</b> | <b>74.5</b> | <b>76.6</b> | <b>79.0</b> | <b>76.6</b> | <b>77.5</b> | <b>78.9</b> | <b>80.3</b> |  |
| ( $\Delta$ )         |              |           | (+2.3)      | (+1.8)      | (+1.7)      | (+1.5)      | (+2.1)      | (+1.9)      | (+1.4)      | (+1.3)      |  |
| U2PL [50]            | [CVPR'22]    |           | 69.0        | 73.0        | 76.3        | 78.6        | 74.9        | 76.5        | 78.5        | 79.1        |  |
| U2PL [50] + MPMC     |              |           | <b>71.2</b> | <b>74.6</b> | <b>77.8</b> | <b>79.8</b> | <b>77.2</b> | <b>78.2</b> | <b>80.0</b> | <b>80.0</b> |  |
| ( $\Delta$ )         |              |           | (+2.2)      | (+1.6)      | (+1.5)      | (+1.2)      | (+2.3)      | (+1.7)      | (+1.5)      | (+0.9)      |  |
| UniMatch [57]        | [CVPR'23]    |           | 75.0        | 76.8        | 77.5        | 78.6        | 76.6        | 77.9        | 79.2        | 79.5        |  |
| UniMatch [57] + MPMC |              |           | <b>76.8</b> | <b>78.0</b> | <b>78.6</b> | <b>79.5</b> | <b>78.5</b> | <b>79.2</b> | <b>80.9</b> | <b>81.3</b> |  |
| ( $\Delta$ )         |              |           | (+1.8)      | (+1.2)      | (+1.1)      | (+0.9)      | (+1.9)      | (+1.3)      | (+1.7)      | (+1.8)      |  |

**Table 3:** Quantitative results of different semi-supervised segmentation methods on ACDC dataset with 3/7 labeled cases. We report Dice Similarity Coefficient (DSC) metric averaged over 3 classes. We use ResNet-101 for all experiments in this dataset.

| Method               |              | 3 cases     | 7 cases     |
|----------------------|--------------|-------------|-------------|
| <i>Supervised</i>    |              | 41.5        | 62.5        |
| UA-MT [60]           | [MICCAI'19]  | 61.0        | 81.5        |
| CPS [7]              | [CVPR'21]    | 60.3        | 83.3        |
| CNN & Trans [34]     | [PMLR'22]    | 65.6        | 86.4        |
| AugSeg [65]          | [CVPR'21]    | 85.1        | 87.3        |
| AugSeg [65] + MPMC   |              | <b>88.6</b> | <b>88.4</b> |
| ( $\Delta$ )         |              | (+3.5)      | (+1.1)      |
| AEL [20]             | [NeurIPS'21] | 82.7        | 87.6        |
| AEL [20] + MPMC      |              | <b>85.4</b> | <b>88.5</b> |
| ( $\Delta$ )         |              | (+2.7)      | (+0.9)      |
| U2PL [50]            | [CVPR'22]    | 84.2        | 89.2        |
| U2PL [50] + MPMC     |              | <b>87.4</b> | <b>90.5</b> |
| ( $\Delta$ )         |              | (+3.2)      | (+1.3)      |
| UniMatch [57]        | [CVPR'23]    | 88.9        | 89.9        |
| UniMatch [57] + MPMC |              | <b>91.8</b> | <b>90.4</b> |
| ( $\Delta$ )         |              | (+2.9)      | (+0.5)      |


### 4.3 Ablation Studies

We conduct experiments to study the impact of various components of our approach. All experiments in this section were conducted on  $\frac{1}{16}$  data partition of the Cityscapes dataset, and use AugSeg [65] as the baseline.


**Segmentation performance at different scales.** Here we analyse how scaling of input features for patch-classification, improves the performance of the SSS method. We consider two scales - average pool kernels  $7 \times 7$  and  $19 \times 19$ . Fig. 4 illustrates for two random classes pole and rider, that average pooling with a kernel size of  $7 \times 7$  is better at enhancing the pixel accuracy for instances with smaller areas, whereas the  $19 \times 19$  kernel size is better at improving accuracy for instances with larger areas. This can be attributed to the larger receptive field of the  $19 \times 19$  kernel, enabling more effective contextual learning for larger regions. While the smaller receptive field of  $7 \times 7$  kernel is better able to discern smaller objects.

**Analysing class features** The crux of our method is to better discriminate pixel regions between different classes. Here we analyze if MPMC integrated to AugSeg [65] can better discriminate features than the baseline for different classes in Cityscapes Dataset. We use the pixel features from the penultimate layer of the segmentation network to analyse via t-SNE. Fig 5 illustrates for four random classes - car, bus, train and truck, integration of MPMC leads to better discrimination between the classes, *i.e.*, compared to AugSeg [65], the features of *train* and *truck* are more clustered and separated from the rest.

**Influence of MPMC on patch-level segmentation.** Our method MPMC improves segmentation by introducing patch classification as an auxiliary task. We analyse at patch level the correlation between multi-label accuracy and the Mean IoU performance of the segmentation network. We use (1 - hamming loss) [48] as a measure of multi-label classification and use the basic teacher-



**Fig. 4: Pixel accuracy across different scales with respect to instance size.** For two different scales, we compare the pixel accuracy with respect to size of instances for two random classes: pole and rider.



**Fig. 5: Analysis of class features** (a) t-SNE of features predicted by AugSeg [65], (b) t-SNE of features predicted by AugSeg [65] + MPMC.

student framework mentioned in Section 3.1 as our baseline. Table 4 illustrates that the avg MIoU improvement is highest (4.05%) in patches having more than 90% multi-label classification accuracy showing that MPMC improves segmentation at the patch level.

#### Impact of MPMC on semantic segmentation in limited labeled data scenarios.

In Table 5 shows the impact of MPMC on segmentation when trained with only labeled data. We use the basic teacher-student framework mentioned in Section 3.1 as our baseline, which achieves MIoU of 66.3, 72.8 and 75.0 in 186, 372 and 744 partitions of labeled images respectively. After introducing PMC, there is an improvement of 1.9%, 1.5% and 1.3% in 186, 372 and 744 partitions of labeled images respectively. We observe that MPMC improves segmentation performance in limited-labeled data scenarios, with the improvement highest for least labeled data, i.e., 1.9% for 186 labeled images.

**Table 4:** Analysis of the influence of MPMC on segmentation at patch-level.

| 1 - hamming loss  | 90+   | 80-90 | 70-80 | 60-70 | 50-60 | 40-50 | 30-40 | 20-30 | 10-20 | 0-10  |
|-------------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| % patches         | 12.6  | 25.2  | 20.5  | 8.4   | 5.1   | 5.2   | 4.1   | 10.2  | 2.4   | 6.3   |
| Avg MIOU $\Delta$ | +4.05 | +1.57 | +0.88 | +1.55 | +0.91 | +1.15 | +0.49 | +0.10 | -1.90 | -2.20 |

**Table 5:** Impact of MPMC on segmentation (MIOU) when trained with different volumes of labeled data. MPMC shows the highest improvement (1.9%) in the least labeled data partition.

|                | # Labeled Data |             |             |
|----------------|----------------|-------------|-------------|
|                | 186            | 372         | 744         |
| w/o MPMC       | 66.3           | 72.8        | 75.0        |
| w/ MPMC        | <b>68.2</b>    | <b>74.3</b> | <b>76.3</b> |
| Gain/ $\Delta$ | +1.9           | +1.5        | +1.3        |

**Table 6:** Ablation study of different component: MPMC on labeled images (MPMCL) and MPMC on unlabeled images (MPMCUL)

| MPMCL | MPMCUL | 1/16        | 1/8         | 1/4         |
|-------|--------|-------------|-------------|-------------|
|       |        | (186)       | (372)       | (744)       |
| ✓     |        | 68.4        | 73.5        | 74.2        |
|       |        | 70.8        | 75.4        | 75.7        |
|       | ✓      | 71.3        | 76.0        | 76.6        |
| ✓     | ✓      | <b>73.1</b> | <b>76.7</b> | <b>77.3</b> |

**Analysis of different components** We ablate each component of MPMC with basic teacher-student framework (Section 3.1) as the baseline. MPMC on labeled images (MPMCL), which introduces multi-scale patch-based classification on only the labeled data, achieves an improvement of 2.4%, 1.9%, and 1.5% under  $\frac{1}{16}$ ,  $\frac{1}{8}$  and  $\frac{1}{4}$  partition protocols respectively. MPMC on unlabeled images (MPMCUL) leads to a further improvement of 2.9%, 2.5%, and 2.4% under  $\frac{1}{16}$ ,  $\frac{1}{8}$  and  $\frac{1}{4}$  partition protocols respectively. Thus, MPMC is effective on both labeled and unlabeled instances in SSS pipelines.

## 5 Conclusion

In this paper, we have proposed a novel Multi-scale Patch-based Multi-label Classifier (MPMC), which introduces patch-level context information in the form of an auxiliary task of multi-scale patch-based multi-label classification. When jointly trained with a SSS model, the patch classifier provides patch-level labels, identifying the classes present within a region. This facilitates the elimination of distractors and enhances the classification of small object segments by the baseline method. Further, the MPMC predictions are used to learn an adaptive weight that reduces the influence of noisy pseudo-labels in the teacher-student framework. We performed a thorough analysis of the importance of multi-scale patch-level contextual information in improving segmentation especially in low data scenarios. Further, we demonstrated the efficacy of our proposed parameterized module MPMC by integrating it into four SSS methodologies and showing consistent improvements in Cityscapes, Pascal VOC and ACDC benchmarks.

**Acknowledgement.** This work was partially supported by the National Science Foundation (IIS-2245652, IIS-2123920, IIS-2212046).

## References

1. Alonso, I., Sabater, A., Ferstl, D., Montesano, L., Murillo, A.C.: Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8219–8228 (2021)
2. Bachman, P., Alsharif, O., Precup, D.: Learning with pseudo-ensembles. *Advances in neural information processing systems* **27** (2014)
3. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging* **37**(11), 2514–2525 (2018)
4. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems* **32** (2019)
5. Chen, H., Jin, Y., Jin, G., Zhu, C., Chen, E.: Semisupervised semantic segmentation by improving prediction confidence. *IEEE Transactions on Neural Networks and Learning Systems* (2021)
6. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 801–818 (2018)
7. Chen, X., Yuan, Y., Zeng, G., Wang, J.: Semi-supervised semantic segmentation with cross pseudo supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2613–2622 (2021)
8. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
9. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552* (2017)
10. Durasov, N., Dorndorf, N., Le, H., Fua, P.: Zigzag: Universal sampling-free uncertainty estimation through two-step inference. *Transactions on Machine Learning Research* (2024)
11. Durasov, N., Oner, D., Donier, J., Le, H., Fua, P.: Enabling uncertainty estimation in iterative neural networks. In: Forty-first International Conference on Machine Learning (2024)
12. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International journal of computer vision* **88**(2), 303–338 (2010)
13. Fan, J., Gao, B., Jin, H., Jiang, L.: Ucc: Uncertainty guided cross-head co-training for semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9947–9956 (2022)
14. French, G., Laine, S., Aila, T., Mackiewicz, M., Finlayson, G.: Semi-supervised semantic segmentation needs strong, varied perturbations. *arXiv preprint arXiv:1906.01916* (2019)
15. Ge, W., Yang, S., Yu, Y.: Multi-evidence filtering and fusion for multi-label classification, object detection and semantic segmentation based on weakly supervised learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1277–1286 (2018)



16. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. *Advances in neural information processing systems* **17** (2004)
17. Guo, H., Wang, S.: Long-tailed multi-label visual recognition by collaborative training on uniform and re-balanced samplings. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15089–15098 (2021)
18. Hariharan, B., Arbeláez, P., Bourdev, L., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: *2011 international conference on computer vision*. pp. 991–998. IEEE (2011)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
20. Hu, H., Wei, F., Hu, H., Ye, Q., Cui, J., Wang, L.: Semi-supervised semantic segmentation via adaptive equalization learning. *Advances in Neural Information Processing Systems* **34**, 22106–22118 (2021)
21. Huang, J., Li, G., Huang, Q., Wu, X.: Learning label specific features for multi-label classification. In: *2015 IEEE International Conference on Data Mining*. pp. 181–190. IEEE (2015)
22. Huo, X., Xie, L., He, J., Yang, Z., Zhou, W., Li, H., Tian, Q.: Atso: Asynchronous teacher-student optimization for semi-supervised image segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1235–1244 (2021)
23. Ke, Z., Qiu, D., Li, K., Yan, Q., Lau, R.W.: Guided collaborative training for pixel-wise semi-supervised learning. In: *European conference on computer vision*. pp. 429–445. Springer (2020)
24. Le, H., Goncalves, B., Samaras, D., Lynch, H.: Weakly labeling the antarctic: The penguin colony case. In: *CVPRW* (June 2019)
25. Le, H., Samaras, D., Lynch, H.J.: A convolutional neural network architecture designed for the automated survey of seabird colonies. *Remote Sensing in Ecology and Conservation* **8**(2), 251–262 (2022)
26. Le, H., Vicente, T.F.Y., Nguyen, V., Hoai, M., Samaras, D.: A+D Net: Training a shadow detector with adversarial shadow attenuation. In: *European Conference on Computer Vision(ECCV)* (2018)
27. Le, H., Yu, C.P., Zelinsky, G., Samaras, D.: Co-localization with category-consistent features and geodesic distance propagation. In: *ICCVW* (2017)
28. Lee, D.H., et al.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: *Workshop on challenges in representation learning, ICML*. vol. 3, p. 896 (2013)
29. Li, S., He, Y., Zhang, W., Zhang, W., Tan, X., Han, J., Ding, E., Wang, J.: Cfcg: Semi-supervised semantic segmentation via cross-fusion and contour guidance supervision. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 16348–16358 (2023)
30. Lin, D.: Probability guided loss for long-tailed multi-label image classification. In: *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence. AAAI’23/IAAI’23/EAAI’23*, AAAI Press (2023). <https://doi.org/10.1609/aaai.v37i2.25244> <https://doi.org/10.1609/aaai.v37i2.25244>
31. Lin, G., Shen, C., Van Den Hengel, A., Reid, I.: Efficient piecewise training of deep structured models for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3194–3203 (2016)

32. Liu, W., Rabinovich, A., Berg, A.C.: Parsenet: Looking wider to see better. arXiv preprint arXiv:1506.04579 (2015)
33. Liu, Y., Tian, Y., Chen, Y., Liu, F., Belagiannis, V., Carneiro, G.: Perturbed and strict mean teachers for semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4258–4267 (2022)
34. Luo, X., Hu, M., Song, T., Wang, G., Zhang, S.: Semi-supervised medical image segmentation via cross teaching between cnn and transformer. In: International Conference on Medical Imaging with Deep Learning. pp. 820–833. PMLR (2022)
35. McClosky, D., Charniak, E., Johnson, M.: Effective self-training for parsing. In: Proceedings of the Human Language Technology Conference of the NAACL, Main Conference. pp. 152–159 (2006)
36. Oliver, A., Odena, A., Raffel, C.A., Cubuk, E.D., Goodfellow, I.: Realistic evaluation of deep semi-supervised learning algorithms. *Advances in neural information processing systems* **31** (2018)
37. Olsson, V., Tranheden, W., Pinto, J., Svensson, L.: Classmix: Segmentation-based data augmentation for semi-supervised learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1369–1378 (2021)
38. Ouali, Y., Hudelot, C., Tami, M.: Semi-supervised semantic segmentation with cross-consistency training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12674–12684 (2020)
39. Read, J., Pfahringer, B., Holmes, G., Frank, E.: Classifier chains for multi-label classification. *Machine learning* **85**, 333–359 (2011)
40. Reif, S., Seibold, C., Freytag, A., Rodner, E., Stiefelhagen, R.: Every annotation counts: Multi-label deep supervision for medical image segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9532–9542 (2021)
41. Ridnik, T., Ben-Baruch, E., Zamir, N., Noy, A., Friedman, I., Protter, M., Zelnik-Manor, L.: Asymmetric loss for multi-label classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 82–91 (2021)
42. Rizve, M.N., Duarte, K., Rawat, Y.S., Shah, M.: In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. arXiv preprint arXiv:2101.06329 (2021)
43. Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in neural information processing systems* **29** (2016)
44. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Gradcam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017)
45. Shi, W., Gong, Y., Ding, C., Tao, Z.M., Zheng, N.: Transductive semi-supervised deep learning using min-max features. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 299–315 (2018)
46. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems* **33**, 596–608 (2020)
47. Sun, R., Mai, H., Zhang, T., Wu, F.: DAW: Exploring the better weighting function for semi-supervised semantic segmentation. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=KR1G7NJUCD>

48. Tsoumakas, G., Katakis, I.: Multi-label classification. *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications: Concepts, Methodologies, Tools, and Applications* **3**, 64 (2008)
49. Verma, V., Kawaguchi, K., Lamb, A., Kannala, J., Bengio, Y., Lopez-Paz, D.: Interpolation consistency training for semi-supervised learning. *arXiv preprint arXiv:1903.03825* (2019)
50. Wang, Y., Wang, H., Shen, Y., Fei, J., Li, W., Jin, G., Wu, L., Zhao, R., Le, X.: Semi-supervised semantic segmentation using unreliable pseudo-labels. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4248–4257 (2022)
51. Wu, T., Huang, Q., Liu, Z., Wang, Y., Lin, D.: Distribution-balanced loss for multi-label classification in long-tailed datasets. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV* 16. pp. 162–178. Springer (2020)
52. Xu, H.M., Liu, L., Bian, Q., Yang, Z.: Semi-supervised semantic segmentation with prototype-based consistency regularization. *arXiv preprint arXiv:2210.04388* (2022)
53. Xu, J., Le, H.: Generating representative samples for few-shot classification. In: *CVPR* (2022)
54. Xu, J., Le, H., Huang, M., Athar, S., Samaras, D.: Variational feature disentangling for fine-grained few-shot classification. In: *ICCV* (2021)
55. Xu, J., Le, H., Samaras, D.: Generating features with increased crop-related diversity for few-shot object detection. In: *CVPR* (2023)
56. Xu, J., Le, H.M., Nguyen, V., Ranjan, V., Samaras, D.: Zero-shot object counting. *CVPR* pp. 15548–15557 (2023)
57. Yang, L., Qi, L., Feng, L., Zhang, W., Shi, Y.: Revisiting weak-to-strong consistency in semi-supervised semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7236–7246 (2023)
58. Yang, L., Zhuo, W., Qi, L., Shi, Y., Gao, Y.: St++: Make self-training work better for semi-supervised semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4268–4277 (2022)
59. Yarowsky, D.: Unsupervised word sense disambiguation rivaling supervised methods. In: *33rd annual meeting of the association for computational linguistics*. pp. 189–196 (1995)
60. Yu, L., Wang, S., Li, X., Fu, C.W., Heng, P.A.: Uncertainty-aware self-ensembling model for semi-supervised 3d left atrium segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 605–613. Springer (2019)
61. Yu, Y., Pedrycz, W., Miao, D.: Multi-label classification by exploiting label correlations. *Expert Systems with Applications* **41**(6), 2989–3004 (2014)
62. Yuan, J., Liu, Y., Shen, C., Wang, Z., Li, H.: A simple baseline for semi-supervised semantic segmentation with strong data augmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8229–8238 (2021)
63. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6023–6032 (2019)
64. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412* (2017)
65. Zhao, Z., Yang, L., Long, S., Pi, J., Zhou, L., Wang, J.: Augmentation matters: A simple-yet-effective approach to semi-supervised semantic segmentation. In: *Pro-*

- ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11350–11359 (2023)
66. Zhong, Y., Yuan, B., Wu, H., Yuan, Z., Peng, J., Wang, Y.X.: Pixel contrastive-consistent semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7273–7282 (2021)
  67. Zhou, T., Wang, S., Bilmes, J.: Time-consistent self-supervision for semi-supervised learning. In: International Conference on Machine Learning. pp. 11523–11533. PMLR (2020)
  68. Zhou, Y., Sun, X., Zha, Z.J., Zeng, W.: Context-reinforced semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4046–4055 (2019)
  69. Zou, Y., Zhang, Z., Zhang, H., Li, C.L., Bian, X., Huang, J.B., Pfister, T.: Pseudoseg: Designing pseudo labels for semantic segmentation. arXiv preprint arXiv:2010.09713 (2020)
  70. Zuo, S., Yu, Y., Liang, C., Jiang, H., Er, S., Zhang, C., Zhao, T., Zha, H.: Self-training with differentiable teacher. arXiv preprint arXiv:2109.07049 (2021)