

# Controller-Guided Partial Label Consistency Regularization with Unlabeled Data

Qian-Wei Wang<sup>1,2</sup>, Bowen Zhao<sup>1</sup>, Mingyan Zhu<sup>1</sup>, Tianxiang Li<sup>1</sup>, Zimo Liu<sup>2</sup>, Shu-Tao Xia<sup>1,2\*</sup>

<sup>1</sup>Tsinghua Shenzhen International Graduate School, Tsinghua University

<sup>2</sup>Research Center of Artificial Intelligence, Peng Cheng Laboratory

{wanggw21, zbm18, zmy20, litx21}@mails.tsinghua.edu.cn, liuzm@pcl.ac.cn, xiast@sz.tsinghua.edu.cn

## Abstract

Partial label learning (PLL) learns from training examples each associated with multiple candidate labels, among which only one is valid. In recent years, benefiting from the strong capability of dealing with ambiguous supervision and the impetus of modern data augmentation methods, consistency regularization-based PLL methods have achieved a series of successes and become mainstream. However, as the partial annotation becomes insufficient, their performances drop significantly. In this paper, we leverage easily accessible unlabeled examples to facilitate the partial label consistency regularization. In addition to a partial supervised loss, our method performs a controller-guided consistency regularization at both the label-level and representation-level with the help of unlabeled data. To minimize the disadvantages of insufficient capabilities of the initial supervised model, we use the controller to estimate the confidence of each current prediction to guide the subsequent consistency regularization. Furthermore, we dynamically adjust the confidence thresholds so that the number of samples of each class participating in consistency regularization remains roughly equal to alleviate the problem of class-imbalance. Experiments show that our method achieves satisfactory performances in more practical situations, and its modules can be applied to existing PLL methods to enhance their capabilities.

## Introduction

In real-world applications, data with unique and correct label is often too costly to obtain (Zhou 2018; Li, Guo, and Zhou 2019; Wang, Yang, and Li 2020; Liu et al. 2023; Chen et al. 2023). Instead, users with varying knowledge and cultural backgrounds tend to annotate the same image with different labels. Traditional supervised learning framework base on "one instance one label" assumption is out of its capability when faced with such ambiguous examples, while partial label learning (PLL) provides an effective solution. Conceptually speaking, PLL (Hüllermeier and Beringer 2006; Nguyen and Caruana 2008; Cour, Sapp, and Taskar 2011; Zhang and Yu 2015; Yu and Zhang 2017; Feng et al. 2020; Lyu et al. 2021; Lv et al. 2020; Wen et al. 2021; Wang et al. 2023; Shi et al. 2023) learns from ambiguous labeling information where each training example is associated with a

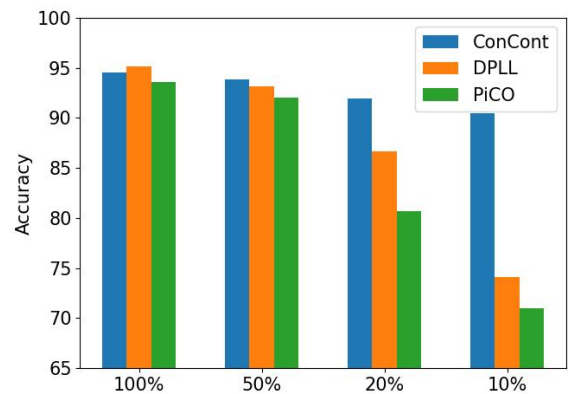


Figure 1: The performances of ConCont (with additional unlabeled examples), DPLL and PiCO with 100%, 50%, 20% and 10% partially annotated examples on CIFAR-10 dataset.

set of candidate labels, among which only one is assumed to be valid. The key to accomplish this task is to find the correct correspondence between each training example and its ground-truth label from the ambiguous candidate collections, i.e. disambiguation. In recent years, partial label learning manifested its capability in solving a wide range of applications such as multimedia contents analysis (Zeng et al. 2013; Chen, Patel, and Chellappa 2018; Xie and Huang 2018), web mining (Jie and Orabona 2010), ecoinformatics (Liu and Dietterich 2012), etc.

With the explosive research of PLL in the deep learning paradigm, consistency regularized disambiguation-based methods (Wang et al. 2022b; Wu, Wang, and Zhang 2022; Wang et al. 2022c; Li et al. 2023; Xia et al. 2023) have achieved significantly better results than other solutions, and have gradually become the mainstream. Such methods usually perturb the samples in the feature space without changing their label semantics, and then use various methods in the label space or representation space to make the outputs of different variants consistent. Modern data augmentation methods (Cubuk et al. 2019a,b; DeVries and Taylor 2017) further improve their performances.

However, their performances are achieved when the number of partially labeled examples is sufficient. In real-world

\*Corresponding author.

applications, PLL is usually in a scenario where the labeling resources are constrained, and the adequacy of partial annotation is not guaranteed. Under this circumstance, existing consistency regularization-based methods often fails to achieve satisfactory performances. As shown in Figure 1, DPLL (Wu, Wang, and Zhang 2022) and PiCO (Wang et al. 2022b) achieve state-of-the-art performances when using complete partial training set, but as the proportion of partial example decreases, their accuracies drop significantly. The reason behind this phenomenon is that when the number of labels is scarce and inherently ambiguous, there is not enough supervision information to guide the initial supervised learning of the model, which leads to the convergence of the consistency regularization to the wrong direction and the emergence of problems such as overfitting and class-imbalance.

Witnessing the enormous power of unlabeled examples via consistency regularization (Sohn et al. 2020; Berthelot et al. 2020; Xie et al. 2020), we hope to facilitate partial label consistency regularization through these readily available data. To this end, an effective framework needs to be designed to maximize the potential of partial and unlabeled examples, as well as reasonable mechanisms to guide the model when supervision information is scarce and ambiguous. In this paper, we propose consistency regularization with controller (abbreviated as ConCont). Our method learns from the supervised information in the training targets (i.e., candidate label sets) via a supervised loss, while performing controller-guided consistency regularization at both label- and representation-levels with the help of unlabeled data. To avoid negative regularization, the controller divides the examples as confident or unconfident according to the prior information and the learning state of the model, and applies different label- and representation-level consistency regularization strategies, respectively. Furthermore, we dynamically adjust the confidence thresholds so that the number of samples of each class participating in consistency regularization remains roughly equal to alleviate the problem of class-imbalance.

## Related Work

Traditional PLL methods can be divided into two categories: averaging-based (Hüllermeier and Beringer 2006; Cour, Sapp, and Taskar 2011; Zhang, Zhou, and Liu 2016) and identification-based (Jin and Ghahramani 2002; Liu and Dietterich 2012; Feng and An 2019; Ni et al. 2021). Averaging-based methods treat all the candidate labels equally, while identification-based methods aim at identifying the ground-truth label directly from candidate label set. With the popularity of deep neural networks, PLL has been increasingly studied in deep learning paradigm. Yao et al. (2020) made the first attempt with an entropy-based regularizer enhancing discrimination. Lv et al. (2020) propose a classifier-consistent risk estimator for partial examples that theoretically converges to the optimal point learned from its fully supervised counterpart under mild condition, as well as an effective method progressively identifying ground-truth labels from the candidate sets. Wen et al. (2021) propose a family of loss functions named leveraged weighted loss taking the

trade-offs between losses on partial labels and non-partial labels into consideration, advancing the former method to a more generalized case. Xu et al. (2021) consider the learning case where the candidate labels are generated in an instance-dependent manner.

Recently, consistency regularization-based PLL methods have achieved impressive results, among which two representatives are: PiCO (Wang et al. 2022b) and DPLL (Wu, Wang, and Zhang 2022). They can be seen as performing consistency regularization at the representation-level and label-level, respectively. To be specific, PiCO aligns the representations of the augmented variants of samples belonging to the same class, and calculates a representation prototype for each class, and then disambiguates the label distribution of the sample according to the distance between the sample representation and each class prototype, forming an iterative EM-like optimization. While DPLL aligns the output label distributions of multiple augmented variants to a conformal distribution, which serves as a comprehensiveness of the label distribution for all augmentations. Despite achieving state-of-the-art performances under fully PLL datasets, their consistency regularization rely heavily on the sufficiency of partial annotations, which greatly limits their applications.

Our work is also related with semi-supervised PLL (Wang, Li, and Zhou 2019; Wang and Zhang 2020). Despite similar learning scenarios, previous semi-supervised PLL methods are all based on nearest-neighbor or linear classifiers, and have not been integrated with modern consistency regularized deep learning, which are very different with our method in terms of algorithm implementation.

## Methodology

### Notations

Let  $\mathcal{X} \subseteq \mathbb{R}^d$  be the input feature space and  $\mathcal{Y} = \{1, 2, \dots, C\}$  denote the label space. We attempt to induce a multi-class classifier  $f : \mathcal{X} \mapsto [0, 1]^C$  from partial label training set  $\mathcal{D}_p = \{(\mathbf{x}^i, \mathbf{y}^i) | 1 \leq i \leq p\}$  and an additional unlabeled set  $\mathcal{D}_u = \{\mathbf{x}^i | p+1 \leq i \leq p+u\}$ . Here,  $\mathbf{y}^i = (y_1^i, y_2^i, \dots, y_C^i) \in \{0, 1\}^C$  is the partial label, in which  $y_j^i = 1$  indicates the  $j$ -th label belonging to the candidate set of sample  $\mathbf{x}^i$  and vice versa. Following the basic assumption of PLL, the ground-truth label  $\ell \in \mathcal{Y}$  is inaccessible to the model while satisfies  $y_\ell^i = 1$ . In order to facilitate the unified operation of partial and unlabeled examples, the target vector of unlabeled examples is represented as  $(1, 1, \dots, 1)$ , i.e., the candidate set equals to  $\mathcal{Y}$ , containing no label information. For the classifier  $f$ , we use  $f_j(\mathbf{x})$  to denote the output of classifier  $f$  on label  $j$  given input  $\mathbf{x}$ .

### Consistency Regularization with Controller

Briefly, our method learns from the supervised information in the training targets (i.e., candidate label sets) via a supervised loss, while performing controller-guided consistency regularization at both label- and representation-level with the help of unlabeled data. The consistency regularization here works by aligning the outputs of different augmented variants of each example. To prevent the model from falling

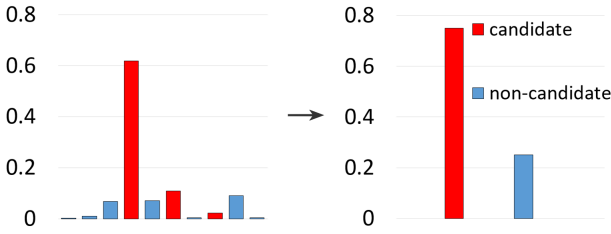


Figure 2: Partial cross-entropy loss transforms multi-class predicted distribution into binary.

into poor convergence due to ambiguity and sparsity of supervision, we design a controller to divide the examples according to the prior information and the learning state of the model to guide the subsequent consistency regularization.

Due to the inaccessibility of unique-label, it is infeasible to train the neural network by minimizing the cross-entropy loss between model prediction and training target as usual. To overcome this difficulty, the original multi-class classification is transformed into a binary classification in our method according to the implication of partial label, i.e. the input should belong to one of the candidate classes. In detail, our method aggregates the predicted probabilities over candidate labels and non-candidate labels respectively (see Figure 2), forming a binary distribution over class *candidate* and class *non-candidate*. Then, the instance is classified to class *candidate* with a simple binary cross-entropy loss.

Given instance  $x$  and its training target  $y$ , the partial cross-entropy loss can be expressed as:

$$\mathcal{L}_{\text{part}} = -\log \sum_{j=1}^C y_j f_j(x). \quad (1)$$

We encourage the model’s output to remain invariant to perturbations applied to input images that do not change the label semantics. With the guidance of controller, our method minimizes the divergence of the output from a pair of different data augmentations:  $(\text{Aug}_w(\cdot), \text{Aug}_s(\cdot))$  at the label-level and  $(\text{Aug}_s(\cdot), \text{Aug}'_s(\cdot))$  at the representation-level (shown in Figure 3), where  $\text{Aug}_w(\cdot)$  denotes the weak augmentation,  $\text{Aug}_s(\cdot)$  and  $\text{Aug}'_s(\cdot)$  denote two different strong augmentations. The reason for this design is that the predicted label distribution from the weakly-augmented variant is more reliable, which is then used to generate the pseudo-label of the input example for subsequent consistency regularization; while previous studies (He et al. 2020; Chen et al. 2020) have shown that the representation alignment of more diverse augmented variants is more conducive to improve the model’s ability. Moreover, we also explored the “two-branch” framework, i.e., performs label- and representation-level consistency regularization on the same pair of augmentations  $(\text{Aug}_w(\cdot), \text{Aug}_s(\cdot))$ . However, the representation learning of this approach is less effective, meanwhile, sharing the same pair of augmentations may lead to overfitting. Experiments also confirmed that its performance is not as good as the adopted “three-branch” version.

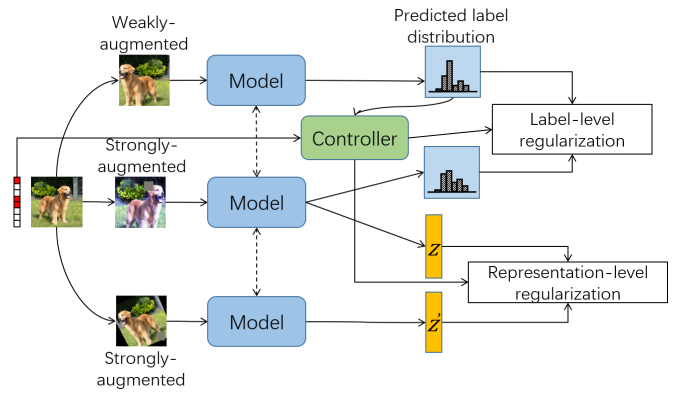


Figure 3: The overall framework of ConCont. Dotted lines represent sharing backbone.

Let  $\tilde{y}^i = f(\text{Aug}_w(x^i))$  be the predicted label distribution of the weakly-augmented variant, we predict the pseudo-label of example  $(x^i, y^i)$  by selecting the candidate class with the maximum predicted probability:  $\hat{y}^i = \arg \max(\tilde{y}^i \circ y^i)$ , where  $\circ$  means the vector point-wise multiplication. To estimate how confident the model is about the example, the controller computes a partial confidence score (p-score for short)  $s^i$  and based on which classifies it as confident or unconfident yet. Given the partial label  $\hat{y}$  and predicted label distribution  $\tilde{y}$ , the p-score for example  $(x, y)$  consists of the following three terms (for the sake of simplicity, we omit the superscript  $i$ ):

- Label information: The prior label information from the provided training target. The fewer the number of candidate classes, the greater the amount of information and vice versa. This term is calculated as  $p_1 = \frac{1}{\sum_{j=1}^C y_j}$ .
- Candidate margin: Measure the prominence in probability between pseudo-label and other candidate classes. We think that the probability margin is more suitable for measuring the confidence than simply the predicted probability over the class. For example, for an instance with 5 candidate classes, the model is obviously more confident in  $\tilde{y} = (0.5, 0.2, 0.1, 0.1, 0.1)$ , comparing with  $\tilde{y}' = (0.5, 0.49, 0.01, 0.0, 0.0)$ . This term is calculated as  $p_2 = \frac{\max(\tilde{y} \circ y, 1) - \max(\tilde{y} \circ y, 2)}{\sum_{j=1}^C \tilde{y}_j y_j}$ , where  $\max(\cdot, k)$  returns the  $k$ -th largest value of the input vector.
- Supervised learning state: The average predicted probability over non-candidate classes, measuring the current learning state of the model from partial supervision, written as  $p_3 = \frac{1 - \sum_{j=1}^C \tilde{y}_j y_j}{C - \sum_{j=1}^C y_j}$ .

The overall p-score is computed as the direct sum of the three terms.

$$p = p_1 + p_2 + p_3. \quad (2)$$

The Precision-Recall (P-R) curves of classifying confident or unconfident examples using p-scores and traditional maximum probabilities are shown in Figure 4. It can be seen that although the curve of p-score has some fluctuations when recall is small, when recall is greater than 60%

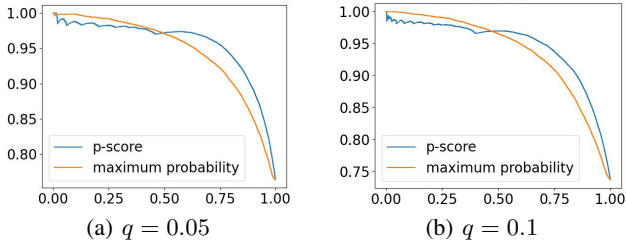


Figure 4: The P-R curves of classifying confident or unconfident examples using p-scores and maximum probabilities, experiments conducted on CIFAR-100 (40%).

(in most cases), the precision of p-score is stably higher than the traditional maximum probability.

The label-level regularization is achieved by aligning the output label distributions of strongly-augmented variants with their pseudo-labels for confident samples with p-scores above the threshold. The loss function of example  $(x, y)$  can be written as:

$$\mathcal{L}_{\text{reg}} = -\mathbf{1}(p \geq \tau_{\hat{y}}) \sum_{j=1}^c \hat{y}_j \log f_j(\text{Aug}_s(x)). \quad (3)$$

Here,  $\hat{y}_j$  is the  $j$ -th component of the one-hot vector corresponding to the pseudo-label and  $\mathbf{1}(\cdot)$  returns 1 if the inside predicate holds, or otherwise returns 0.  $\tau_{\hat{y}}$  is the adaptive threshold for class  $\hat{y}$ , which will be explained in detail in the next section.

Supervised contrastive learning (Khosla et al. 2020) drawing the projected representations of different augmented variants from the samples of the same class closer and those from different classes apart, is far more effective than simply minimizing the squared loss between two augmented representations and serves as a good way for representation-level consistency regularization. In this paper, since neither partial label nor unlabeled examples have valid labels, we can only proceed with the help of predicted label distributions from  $\text{Aug}_w(\cdot)$ . The controller selects positive and negative samples for each example  $(x, y)$  (whose predicted representation denoted as  $z$ ) according to the pseudo-labels and p-scores, and the representation distance between positive samples (representation denoted as  $z_+$ ) and  $x$  are shortened, and the distance between negative samples (representation denoted as  $z_-$ ) and  $x$  are lengthened. Let  $P(x)$  and  $N(x)$  separately denote the set of selected positives and negatives, the loss function can be expressed as:

$$\mathcal{L}'_{\text{reg}} = -\frac{1}{|P(x)|} \sum_{z_+ \in P(x)} \log \frac{\exp(z^\top z_+/t)}{\sum_{z_+ \in P(x)} \exp(z^\top z_+/t) + \sum_{z_- \in N(x)} \exp(z^\top z_-/t)}, \quad (4)$$

where  $t \geq 0$  is the temperature.

Then, the problem comes to the selection of positive and negative sample. In order to ensure the accuracy of the positive set, at the same time make the negative set as large as

possible, which is proven to be crucial by previous studies, our selection strategy is as follows: 1. For a confident input, the confident samples with the same class prediction are selected as positives and all other samples with the different prediction are negatives; 2. For an unconfident input, only the representation of another augmented variant  $\text{Aug}'_s(\cdot)$  of it is selected as positive and the same with confident input, all sample with different prediction are negatives. To keep the purity of positive set, when facing a sample with the same predicted class as the input, but we are uncertain about one or both of the predictions, the sample is ignored at the current step, i.e., not selected as positive or negative. Since the proportion of such samples is usually lower than  $1/C$ , discarding them doesn't have a noticeable impact on the total number of contrast representations. In ablation studies, we compare four strategies for sample selection for representation-level consistency regularization and our proposal outperforms the others.

Putting the supervised and the consistency regularized losses together, the overall training objective is:

$$\mathcal{L} = \mathcal{L}_{\text{part}} + \lambda \mathcal{L}_{\text{reg}} + \mu \mathcal{L}'_{\text{reg}}, \quad (5)$$

where  $\lambda$  and  $\mu$  are loss weights.

### Adaptive Confidence Threshold

PLL often faces severe class-imbalance, especially when there is less partially annotated examples. It is worth noting that the class-imbalance encountered here has different causes from the one we often discuss in supervised learning and long-tailed PLL (Wang et al. 2022a; Xu and Lian 2023; Hong et al. 2023), which is usually caused by the long-tail intrinsic of the dataset. In PLL, the main reason for class-imbalance is that the provided ambiguous supervision allows the model to classify partial examples into any category in the candidate set, and in the consistency regularization lacking of enough correct guidance from initial supervised training model, due to the uneven learning progress of different classes, some strong classes may have so many predictions while the weak classes are gradually overwhelmed.

Inspired by the intrinsic of the problem, we expect to make the consistent regularization process fairer for each class, i.e., make the number of predicted pseudo-labels filtered by the controller for each class as equal as possible. To this end, we propose the adaptive confidence threshold that dynamically adjusts the threshold after each training step for each class, lowering the thresholds for weak classes while raising for strong ones. Specifically, we deem the adaptive thresholds for class 1 to  $C$ , denoted as  $(\tau_1, \tau_2, \dots, \tau_C)$ , as model parameters and minimize the cross-entropy between the uniform distribution and the distribution of the numbers of filtered pseudo-labels via gradient descent.

Let  $s_j = \text{Count}_j(\tau_j)$  denote the count number of pseudo-labels of class  $j$  with p-scores above threshold  $\tau_j$  at the current training step, the distribution of the numbers of filtered pseudo-labels can be expressed as:  $P = (s_1/s', s_2/s', \dots, s_C/s')$ , where  $s' = \sum_{j=1}^C s_j$  is the summation of all filtered pseudo-labels, and the uniform distribution is written as:  $U = (1/C, 1/C, \dots, 1/C)$ .

Dataset	Method	$q = 0.1$	$q = 0.3$	$q = 0.5$	$q = 0.7$
CIFAR-10 (20%)	DPLL	88.74±0.06%	87.15±0.17%	86.66±0.16%	63.92±0.45%
	DPLL*	90.10±0.12%	89.18±0.13%	82.61±0.38%	68.64±0.49%
	PiCO	84.37±0.34%	83.33±0.20%	80.67±0.51%	66.00±1.22%
	PiCO*	91.10±0.18%	90.74±0.50%	58.20±3.78%	29.21±2.47%
	SSPL	55.67±0.22%	52.51±0.29%	49.62±0.32%	46.10±0.41%
	PARM	49.12±0.29%	46.04±0.23%	42.84±0.77%	33.13±3.71%
	DPLL+Controller*	<b>92.95±0.03%</b>	<b>92.50±0.06%</b>	90.82±0.10%	88.55±0.15%
	ConCont (partial only)	88.45±0.08%	86.96±0.07%	85.49±0.21%	82.73±0.32%
	ConCont (two-branch)	92.65±0.03%	92.27±0.11%	91.44±0.12%	91.05±0.15%
	ConCont	92.59±0.07%	92.22±0.08%	<b>91.94±0.10%</b>	<b>91.49±0.06%</b>
CIFAR-10 (10%)	DPLL	81.55±0.12%	79.20±0.24%	74.09±0.37%	28.08±0.83%
	DPLL*	85.24±0.13%	80.57±0.37%	71.24±0.36%	15.39±0.80%
	PiCO	77.87±0.44%	76.15±0.30%	71.01±0.68%	51.61±1.69%
	PiCO*	81.96±0.53%	57.19±2.88%	38.28±2.02%	20.02±1.36%
	SSPL	53.17±0.33%	49.75±0.10%	46.79±0.73%	44.10±0.73%
	PARM	50.19±0.74%	47.15±0.14%	42.16±1.18%	30.69±2.05%
	DPLL+Controller*	91.00±0.14%	90.37±0.21%	85.71±0.36%	73.74±0.85%
	ConCont (partial only)	82.31±0.17%	80.57±0.35%	76.69±0.33%	61.77±0.69%
	ConCont (two-branch)	90.83±0.05%	90.39±0.10%	90.23±0.09%	81.34±0.24%
	ConCont	<b>91.23±0.02%</b>	<b>90.89±0.11%</b>	<b>90.46±0.16%</b>	<b>82.20±0.71%</b>
Fully Supervised	94.91±0.07%				

Table 1: Accuracy (mean±std) comparisons on CIFAR-10 (10% and 20%) with different ambiguity levels.

The partial derivative of the cross-entropy loss  $H(U; P)$  with respect to threshold  $\tau_j$  is calculated as:

$$\begin{aligned} \frac{\partial H(U; P)}{\partial \tau_j} &= \frac{\partial H(U; P)}{\partial s_j} \cdot \frac{\partial s_j}{\partial \tau_j} \\ &= \frac{1}{\ln 2} \cdot \frac{s_j - \bar{s}}{s' \cdot s_j} \cdot \frac{\partial s_j}{\partial \tau_j}, \end{aligned} \quad (6)$$

where  $\bar{s} = s'/C$  is the mean value of the number of filtered pseudo-labels. The detailed derivation of Eq. 6 can be found in Appendix.

The last term of Eq. 6 can also be written as  $Count'_j(\tau_j)$ . The  $Count'_j(\cdot)$  function decreases monotonically from the number of all predictions of class  $j$  to 0 as the input threshold increases, and  $Count'_j(\tau_j)$  reflects the sample density where the p-score is in the neighborhood of  $\tau_j$ . Since samples are fed into the model in batches during training and the number of samples of each class in each batch is too small to estimate  $Count'_j(\tau_j)$ . In actual implementation, we think that the density of samples can be considered to be roughly proportional to the number of filtered pseudo-labels  $s_j$ . The update of  $\tau_j$  through gradient descent turns to be:

$$\tau_j = \tau_j - \frac{\bar{s} - s_j}{s'} \cdot \gamma_\tau, \quad (7)$$

where  $\gamma_\tau$  is the learning rate for adaptive thresholds.

Finally, the thresholds are clamped to a preset interval.

$$\tau_j = \min(\max(\tau_j, \tau_l), \tau_u), \quad (8)$$

where  $\tau_l$  and  $\tau_u$  denote the lower and upper bound of adaptive thresholds, respectively.

## Experiments

### Experimental Setup

We evaluate on commonly used benchmarks CIFAR-10 (Krizhevsky, Hinton et al. 2009) and SVHN (Goodfellow

et al. 2013) with 10% and 20% partial examples, and CIFAR-100 (Krizhevsky, Hinton et al. 2009) with 20% and 40% partial examples. To synthesize partial labels, negative labels are randomly flipped to false positive ones with a probability  $q$ . The comparing methods include state-of-the-art PLL methods DPLL (Wu, Wang, and Zhang 2022) and PiCO (Wang et al. 2022b) (\* indicates the version using unlabeled examples), as well as traditional semi-supervised PLL methods SSPL (Wang, Li, and Zhou 2019) and PARM (Wang and Zhang 2020). ConCont (partial only) shows the results of our method without unlabeled examples, and ConCont (two-branch) is the previously mentioned "two-branch" implementation. To examine the versatility of our method, we additionally adapt the controller module (as well as adaptive thresholds) to DPLL, denoted as DPLL+Controller\*. We record the mean and standard deviation based on three independent runs. Additional details of the experiments can be found in the Appendix.

### Main Results

Table 1, 2 and 5 report the main experimental results. As is shown, our method consistently outperforms all previous methods. For example, on CIFAR-10 dataset with 10% partial examples, our method improves upon the best baseline by 5.99%, 10.32% and 16.37% where  $q$  is set to 0.1, 0.3 and 0.5 respectively. We can also observe that our method achieves decent performances despite some degradation in some extreme cases with rare partial annotation and high ambiguity, where previous methods fail to obtain usable results, which extends the practical applicability of PLL. It is also noteworthy that adapting the controller and adaptive thresholds to DPLL achieves satisfactory results, which shows that the controller inspires the capability of integrating unlabeled examples of DPLL, demonstrating the versatility of our components.

Dataset	Method	$q = 0.1$	$q = 0.3$	$q = 0.5$	$q = 0.7$
SVHN (20%)	DPLL	95.84±0.20%	95.00±0.21%	94.30±0.02%	59.50±1.82%
	DPLL*	90.84±0.24%	83.25±0.38%	62.37±2.21%	24.47±1.83%
	PiCO	94.35±0.19%	42.66±1.16%	20.44±0.49%	13.86±0.44%
	PiCO*	71.93±0.35%	30.29±1.06%	16.79±0.94%	11.93±0.59%
	SSPL	75.43±0.09%	73.93±0.15%	71.96±0.12%	69.00±0.32%
	PARM	70.80±0.57%	68.28±0.61%	63.80±2.26%	61.87±0.66%
	DPLL+Controller*	96.98±0.18%	96.47±0.09%	94.88±0.11%	82.39±0.22%
	ConCont (partial only)	96.40±0.09%	96.39±0.09%	<b>95.96</b> ±0.04%	80.51±0.31%
	ConCont (two-branch)	97.05±0.10%	<b>96.98</b> ±0.02%	94.81±0.09%	90.92±0.10%
	ConCont	<b>97.07</b> ±0.14%	96.92±0.12%	<b>95.68</b> ±0.05%	<b>94.25</b> ±0.07%
SVHN (10%)	DPLL	94.64±0.04%	94.37±0.14%	90.30±0.13%	37.79±2.37%
	DPLL*	79.15±0.64%	55.06±1.69%	19.93±1.02%	19.65±0.80%
	PiCO	89.34±0.16%	60.06±1.00%	16.48±0.24%	13.86±0.33%
	PiCO*	37.60±1.98%	34.30±2.09%	14.55±1.34%	13.25±0.81%
	SSPL	73.54±0.04%	71.58±0.14%	69.21±0.34%	65.51±0.42%
	PARM	67.94±0.96%	66.95±1.36%	62.69±2.73%	59.12±0.71%
	DPLL+Controller*	96.33±0.09%	95.40±0.16%	92.89±0.23%	70.02±1.45%
	ConCont (partial only)	95.12±0.08%	95.04±0.06%	94.25±0.12%	87.67±0.09%
	ConCont (two-branch)	<b>96.72</b> ±0.01%	96.19±0.02%	93.76±0.13%	88.59±0.18%
	ConCont	96.62±0.06%	<b>96.40</b> ±0.10%	<b>94.40</b> ±0.07%	<b>88.71</b> ±0.07%
Fully Supervised	97.35%±0.12%				

Table 2: Accuracy (mean±std) comparisons on SVHN (10% and 20%) with different ambiguity levels.

## Ablation Studies

Ablation	CIFAR-10 (10%, $q = 0.5$ )	CIFAR-100 (20%, $q = 0.1$ )
ConCont	<b>90.46</b> ±0.16%	<b>62.85</b> ±0.26%
w/o $\mathcal{L}_{\text{part}}$	44.32±1.51%	18.84±0.87%
w/o $\mathcal{L}_{\text{reg}}$	80.62±0.34%	55.91±0.35%
w/o $\mathcal{L}'_{\text{reg}}$	86.76±0.13%	60.05±0.19%
w/o $\mathcal{L}_{\text{reg}}$ and $\mathcal{L}'_{\text{reg}}$	70.93±0.21%	49.42±0.37%

Table 3: Ablation study on different components of our method.

**Effects of supervised loss and consistency regularization.** We ablate the contributions of the supervised loss and consistency regularization. Specifically, we compare our method with four degenerations: 1. w/o  $\mathcal{L}_{\text{part}}$  substitutes the partial cross-entropy loss with the traditional cross-entropy loss for regular supervised learning; 2. w/o  $\mathcal{L}_{\text{reg}}$  does not perform label-level consistency regularization; 3. w/o  $\mathcal{L}'_{\text{reg}}$  does not perform representation-level consistency regularization; w/o  $\mathcal{L}_{\text{reg}}$  and  $\mathcal{L}'_{\text{reg}}$  does not perform consistency regularization at both levels. Table 3 reports the accuracy comparisons on 10% partially annotated CIFAR-10 dataset with  $q = 0.5$  and 20% partially annotated CIFAR-100 dataset with  $q = 0.1$ . The bad results achieved by w/o  $\mathcal{L}_{\text{part}}$  indicates that a supervised loss that conforms to the partial label semantics is fundamental. It can also be seen that the consistency regularization at both levels contribute significantly the model performance.

**Comparison of different class-imbalance methods.** To demonstrate the superiority of the adaptive confidence threshold, we further experiment with two widespread class-imbalance methods working on pseudo-labels, i.e., Distri-

Method	CIFAR-10 (10%, $q = 0.5$ )	CIFAR-100 (20%, $q = 0.1$ )
Adaptive	<b>90.46</b> ±0.16%	<b>62.85</b> ±0.26%
DA	89.08±0.25%	61.04±0.02%
CPL	88.83±0.14%	46.32±1.21%
Vanilla	86.50±0.21%	59.50±0.28%
Adaptive	90.32±0.12%	<b>62.86</b> ±0.23%
DA (soft)	88.77±0.15%	61.07±0.13%
Vanilla (soft)	87.95±0.01%	60.10±0.25%

Table 4: Accuracy comparisons of different class-imbalance methods. Vanilla means no additional operation alleviating class-imbalance, and "(soft)" means the usage of KL-divergence of label distribution instead of the cross-entropy loss with one-hot pseudo-label in the label-level consistency regularization.

bution Alignment (DA) from ReMixMatch (Berthelot et al. 2020) and Curriculum Pseudo-Labeling (CPL) from FlexMatch (Zhang et al. 2021). It can be seen from Table 4 that the performance improvements brought by alleviating class-imbalance are significant. Our method brings more than 3% accuracy improvements on both datasets compared to not using and is better than other class-imbalance methods (outperforms the best comparing method by 1.38% on CIFAR-10 (10%,  $q = 0.5$ ) and 1.81% on CIFAR-100 (20%,  $q = 0.1$ )).

Additionally, we explore the usage of KL-divergence of label distribution instead of the cross-entropy loss with one-hot pseudo-labels in the label-level consistency regularization, i.e., soft pseudo-labels. The same with one-hot pseudo-labels, we compute p-scores and mask out unconfident examples via adaptive thresholds. As shown in Table 4, our method still achieves the best performances. CPL is not

Dataset	Method	$q = 0.05$	$q = 0.1$	$q = 0.15$	$q = 0.2$
CIFAR-100 (40%)	DPLL	68.07±0.13%	66.51±0.10%	62.37±0.08%	55.54±0.37%
	DPLL*	68.02±0.26%	64.92±0.27%	54.78±0.56%	45.60±2.25%
	PiCO	57.62±0.15%	40.07±0.28%	26.04±0.69%	17.32±0.47%
	PiCO*	64.01±0.35%	40.06±1.82%	21.52±1.22%	18.61±0.56%
	SSPL	22.20±0.10%	20.33±0.13%	18.95±0.06%	17.94±0.04%
	PARM	/	/	/	/
	DPLL+Controller*	<b>70.17±0.20%</b>	68.10±0.23%	64.33±0.18%	58.02±0.37%
	ConCont (partial only)	66.24±0.10%	64.02±0.13%	60.96±0.26%	56.03±0.37%
	ConCont (two-branch)	68.95±0.10%	67.86±0.18%	<b>66.03±0.22%</b>	59.41±0.27%
	ConCont	69.17±0.29%	<b>68.46±0.16%</b>	65.70±0.21%	<b>60.59±0.18%</b>
CIFAR-100 (20%)	DPLL	58.59±0.27%	51.61±0.17%	39.56±0.44%	28.21±0.23%
	DPLL*	57.25±0.28%	41.03±0.49%	29.36±0.72%	24.71±0.45%
	PiCO	49.05±0.64%	27.67±0.95%	14.85±0.40%	9.59±0.83%
	PiCO*	47.17±0.73%	27.10±1.51%	17.70±1.44%	9.65±1.69%
	SSPL	18.73±0.23%	16.67±0.18%	15.61±0.10%	14.24±0.35%
	PARM	/	/	/	/
	DPLL+Controller*	<b>65.20±0.04%</b>	58.02±0.17%	51.89±0.58%	40.42±1.12%
	ConCont (partial only)	56.43±0.12%	51.63±0.04%	44.32±0.08%	38.60±0.28%
	ConCont (two-branch)	63.86±0.05%	62.00±0.27%	58.07±0.14%	47.30±0.16%
	ConCont	63.55±0.10%	<b>62.85±0.26%</b>	<b>58.62±0.02%</b>	<b>49.42±0.31%</b>
Fully Supervised		73.56±0.10%			

Table 5: Accuracy (mean±std) comparisons on CIFAR-100 (20% and 40%) with different ambiguity levels. ”/” means not applicable.

compared here since it is designed for one-hot pseudo-labels.

Strategy	CIFAR-10 (10%, $q = 0.5$ )	CIFAR-100 (20%, $q = 0.1$ )
Ours	<b>90.46±0.16%</b>	<b>62.85±0.26%</b>
HCPN	89.25±0.16%	62.50±0.17%
HCP	89.76±0.30%	62.35±0.27%
SupCon	89.32±0.08%	62.22±0.11%
UnsupCon	89.70±0.01%	61.68±0.31%
w/o $\mathcal{L}'_{\text{reg}}$	86.76±0.13%	60.05±0.19%
$Y_{\text{true}}$	90.66±0.09%	64.50±1.16%

Table 6: Accuracy comparisons among different strategies.

### Different selection strategies of positive and negative samples.

In order to ensure the accuracy of the positive set, at the same time make the negative set with high error-tolerance as large as possible, our method includes a tailor-made sample selection strategy in the representation-level consistency regularization. Here, we compare our method with other strategies including: UnsupCon, SupCon, HCP and HCPN, their specific ways are listed as follows:

- UnsupCon: Traditional unsupervised contrastive learning, selects the representation of  $\text{Aug}'_s(\cdot)$  as the only positive sample.
- SupCon: Directly use pseudo-labels as ground-truth and performs supervised contrastive learning.
- High-Confidence Positive (HCP): Ensure the accuracy of positive set and put all remaining samples into the negative set. For a confident input, select confident samples with the same prediction as positives; while for an uncon-

fidant input, only select the representation of its another view as positive.

- High-Confidence Positive and Negative (HCPN): Ensure the accuracy of both sets. For a confident input, we separately select confident samples with or without the same pseudo-label as positives or negatives, while neglecting the unconfident input when computing  $\mathcal{L}'_{\text{reg}}$ .

In order to reflect the differences more clearly, we use true unique-labels as the supervision for supervised contrastive learning serving as the upper bound (denoted as  $Y_{\text{true}}$ ), and w/o  $\mathcal{L}'_{\text{reg}}$  as the lower bound.

As shown in table 6, our method performs the best among all comparing strategies on both cases. On CIFAR-10 (10%,  $q = 0.5$ ), our method outperforms the second-ranked HCP by 0.7% and is only 0.2% below the upper bound. On CIFAR-100 (20%,  $q = 0.1$ ), HCPN performs the best among other comparing strategies (mainly because that it only uses high-confident samples and ignores uncertain ones), and our method outperforms HCPN by 0.35%.

## Conclusion

In this paper, we observe that the current mainstream, consistent regularization-based PLL methods are flawed when there are not enough partially labeled examples, which greatly limits their practical usability. To address this issue, we introduce cheaply acquired unlabeled examples and devise a method for consistency regularization at the label and representation levels. The highlights of our approach are the p-scores computed by the controller and how they are used to control the subsequent consistency regularization, as well as the simple yet effective class-balancing mechanism. Experiments show that our method can achieve satisfactory results with only a small number of partial examples and unlabeled examples.

## Acknowledgments

This work is supported in part by the National Natural Science Foundation of China under Grant 62171248, Shenzhen Science and Technology Program (JCYJ20220818101012025), and the PCNL KEY project (PCL2023AS6-1).

## References

- Berthelot, D.; Carlini, N.; Cubuk, E. D.; Kurakin, A.; Sohn, K.; Zhang, H.; and Raffel, C. 2020. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *International Conference on Learning Representations*.
- Chen, C.; Patel, V. M.; and Chellappa, R. 2018. Learning from Ambiguously Labeled Face Images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(7): 1653–1667.
- Chen, H.; Shah, A.; Wang, J.; Tao, R.; Wang, Y.; Xie, X.; Sugiyama, M.; Singh, R.; and Raj, B. 2023. Imprecise label learning: A unified framework for learning with various imprecise label configurations. *arXiv preprint arXiv:2305.12715*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Cour, T.; Sapp, B.; and Taskar, B. 2011. Learning from partial labels. *Journal of Machine Learning Research*, 12: 1501–1536.
- Cubuk, E. D.; Zoph, B.; Mane, D.; Vasudevan, V.; and Le, Q. V. 2019a. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 113–123.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2019b. RandAugment: Practical data augmentation with no separate search. *CoRR*, abs/1909.13719.
- DeVries, T.; and Taylor, G. W. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*.
- Feng, L.; and An, B. 2019. Partial Label Learning with Self-Guided Retraining. In *AAAI*, 3542–3549. AAAI Press.
- Feng, L.; Lv, J.; Han, B.; Xu, M.; Niu, G.; Geng, X.; An, B.; and Sugiyama, M. 2020. Provably Consistent Partial-Label Learning. In *Advances in neural information processing systems*.
- Goodfellow, I. J.; Bulatov, Y.; Ibarz, J.; Arnoud, S.; and Shet, V. 2013. Multi-digit number recognition from street view imagery using deep convolutional neural networks. *arXiv preprint arXiv:1312.6082*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- Hong, F.; Yao, J.; Zhou, Z.; Zhang, Y.; and Wang, Y. 2023. Long-Tailed Partial Label Learning via Dynamic Rebalancing. In *ICLR*.
- Hüllermeier, E.; and Beringer, J. 2006. Learning from ambiguously labeled examples. *Intelligent Data Analysis*, 10(5): 419–439.
- Jie, L.; and Orabona, F. 2010. Learning from candidate labeling sets. In *Advances in Neural Information Processing Systems 23*, 1504–1512. Cambridge, MA: MIT Press.
- Jin, R.; and Ghahramani, Z. 2002. Learning with Multiple Labels. In *NeurIPS*, 897–904. MIT Press.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33: 18661–18673.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, X.; Jiang, Y.; Li, C.; Wang, Y.; and Ouyang, J. 2023. Learning with Partial Labels from Semi-supervised Perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37(7), 8666–8674.
- Li, Y.-F.; Guo, L.-Z.; and Zhou, Z.-H. 2019. Towards safe weakly supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 43(1): 334–346.
- Liu, L.; and Dietterich, T. 2012. A conditional multinomial mixture model for superset label learning. In *Advances in Neural Information Processing Systems 25*, 557–565. Cambridge, MA: MIT Press.
- Liu, S.; Cao, Y.; Zhang, Q.; Feng, L.; and An, B. 2023. Consistent Complementary-Label Learning via Order-Preserving Losses. In *International Conference on Artificial Intelligence and Statistics*, 8734–8748. PMLR.
- Lv, J.; Xu, M.; Feng, L.; Niu, G.; Geng, X.; and Sugiyama, M. 2020. Progressive Identification of True Labels for Partial-Label Learning. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 6500–6510. PMLR.
- Lyu, G.; Feng, S.; Wang, T.; Lang, C.; and Li, Y. 2021. GM-PLL: Graph Matching Based Partial Label Learning. *IEEE Trans. Knowl. Data Eng.*, 33(2): 521–535.
- Nguyen, N.; and Caruana, R. 2008. Classification with partial labels. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 381–389. Las Vegas, NV.
- Ni, P.; Zhao, S.-Y.; Dai, Z.-G.; Chen, H.; and Li, C.-P. 2021. Partial label learning via conditional-label-aware disambiguation. *Journal of Computer Science and Technology*, 36(3): 590–605.
- Shi, Y.; Wu, D.-D.; Geng, X.; and Zhang, M.-L. 2023. Robust Representation Learning for Unreliable Partial Label Learning. *arXiv preprint arXiv:2308.16718*.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C.; Cubuk, E. D.; Kurakin, A.; and Li, C. 2020. Fix-Match: Simplifying Semi-Supervised Learning with Consistency and Confidence. In *NeurIPS*.
- Wang, H.; Xia, M.; Li, Y.; Mao, Y.; Feng, L.; Chen, G.; and Zhao, J. 2022a. Solar: Sinkhorn label refinery for imbalanced partial-label learning. *Advances in Neural Information Processing Systems*, 35: 8104–8117.

- Wang, H.; Xiao, R.; Li, Y.; Feng, L.; Niu, G.; Chen, G.; and Zhao, J. 2022b. PiCO: Contrastive Label Disambiguation for Partial Label Learning. *International Conference on Learning Representations*.
- Wang, H.; Xiao, R.; Li, Y.; Feng, L.; Niu, G.; Chen, G.; and Zhao, J. 2022c. PiCO+: Contrastive Label Disambiguation for Robust Partial Label Learning. *arXiv preprint arXiv:2201.08984*.
- Wang, H.; Yang, S.; Lyu, G.; Liu, W.; Hu, T.; Chen, K.; Feng, S.; and Chen, G. 2023. Deep Partial Multi-Label Learning with Graph Disambiguation. *arXiv preprint arXiv:2305.05882*.
- Wang, Q.-W.; Li, Y.-F.; and Zhou, Z.-H. 2019. Partial Label Learning with Unlabeled Data. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 3755–3761.
- Wang, Q.-W.; Yang, L.; and Li, Y.-F. 2020. Learning from weak-label data: A deep forest expedition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 6251–6258.
- Wang, W.; and Zhang, M.-L. 2020. Semi-supervised partial label learning via confidence-rated margin maximization. *Advances in neural information processing systems*, 33: 6982–6993.
- Wen, H.; Cui, J.; Hang, H.; Liu, J.; Wang, Y.; and Lin, Z. 2021. Leveraged Weighted Loss for Partial Label Learning. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, 11091–11100. PMLR.
- Wu, D.-D.; Wang, D.-B.; and Zhang, M.-L. 2022. Revisiting Consistency Regularization for Deep Partial Label Learning. In *International conference on machine learning*.
- Xia, S.; Lv, J.; Xu, N.; Niu, G.; and Geng, X. 2023. Towards Effective Visual Representations for Partial-Label Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15589–15598.
- Xie, M.-K.; and Huang, S.-J. 2018. Partial multi-label learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32(1).
- Xie, Q.; Dai, Z.; Hovy, E.; Luong, T.; and Le, Q. 2020. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33: 6256–6268.
- Xu, M.; and Lian, Z. 2023. Pseudo Labels Regularization for Imbalanced Partial-Label Learning. *arXiv preprint arXiv:2303.03946*.
- Xu, N.; Qiao, C.; Geng, X.; and Zhang, M.-L. 2021. Instance-dependent partial label learning. *Advances in Neural Information Processing Systems*, 34: 27119–27130.
- Yao, Y.; Deng, J.; Chen, X.; Gong, C.; Wu, J.; and Yang, J. 2020. Deep discriminative cnn with temporal ensembling for ambiguously-labeled image classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 12669–12676.
- Yu, F.; and Zhang, M.-L. 2017. Maximum margin partial label learning. *Machine Learning*, 106(4): 573–593.
- Zeng, Z.; Xiao, S.; Jia, K.; Chan, T.-H.; Gao, S.; Xu, D.; and Ma, Y. 2013. Learning by associating ambiguously labeled images. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 708–715. Portland, OR.
- Zhang, B.; Wang, Y.; Hou, W.; Wu, H.; Wang, J.; Okumura, M.; and Shinozaki, T. 2021. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 34: 18408–18419.
- Zhang, M.-L.; and Yu, F. 2015. Solving the partial label learning problem: An instance-based approach. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 4048–4054. Buenos Aires, Argentina.
- Zhang, M.-L.; Zhou, B.-B.; and Liu, X.-Y. 2016. Partial label learning via feature-aware disambiguation. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1335–1344.
- Zhou, Z.-H. 2018. A brief introduction to weakly supervised learning. *National science review*, 5(1): 44–53.