

Unreliable Partial Label Learning with Recursive Separation

Yu Shi, Ning Xu*, Hua Yuan and Xin Geng*

School of Computer Science and Engineering, Southeast University, Nanjing 211189, China

{seushiyu, xning, yuanhua, xgeng}@seu.edu.cn

Abstract

Partial label learning (PLL) is a typical weakly supervised learning problem in which each instance is associated with a candidate label set, and among which only one is true. However, the assumption that the ground-truth label is always among the candidate label set would be unrealistic, as the reliability of the candidate label sets in real-world applications cannot be guaranteed by annotators. Therefore, a generalized PLL named Unreliable Partial Label Learning (UPLL) is proposed, in which the true label may not be in the candidate label set. Due to the challenges posed by unreliable labeling, previous PLL methods will experience a marked decline in performance when applied to UPLL. To address the issue, we propose a two-stage framework named Unreliable Partial Label Learning with Recursive Separation (UPLLRS). In the first stage, the self-adaptive recursive separation strategy is proposed to separate the training set into a reliable subset and an unreliable subset. In the second stage, a disambiguation strategy is employed to progressively identify the ground-truth labels in the reliable subset. Simultaneously, semi-supervised learning methods are adopted to extract valuable information from the unreliable subset. Our method demonstrates state-of-the-art performance as evidenced by experimental results, particularly in situations of high unreliability. Code and supplementary materials are available at <https://github.com/dhiyu/UPLLRS>.

1 Introduction

Partial label learning (PLL) is a typical weakly supervised learning problem where the candidate label set is given for each instance but among which only one is true. Compared with the ordinary supervised learning problem where each instance is associated with only one ground-truth label, partial label learning induces predictive model from ambiguous labels, hence considerably reduces the cost of data annotations. Nowadays, PLL has been extensively employed in the field

of web mining [Luo and Orabona, 2010], multimedia content analysis [Zeng *et al.*, 2013], automatic image annotations [Chen *et al.*, 2018], ecoinformatics [Liu and Dietterich, 2012; Tang and Zhang, 2017], etc.

A variety of methods have been proposed for addressing the PLL problem. The most common strategy to learn from partial labels is disambiguation, where Identification-Based Strategy (IBS) and Average-Based Strategy (ABS) are two main disambiguation strategies. For IBS, iterative optimization is employed to predict true label treated as latent variable. While ABS treats all labels in the candidate label set in an equal manner where probabilities of modeling outputs are averaged to get the final prediction. It memorizes all candidate labels, since it avoids identifying the latent ground truth label. Recently, deep neural network based IBS method have achieved promising performance on PLL. PiCO [Wang *et al.*, 2022] achieves a significant improvement in performance by adopting contrastive learning strategy in PLL, which is able to learn high quality representation. CR-DPLL [Wu *et al.*, 2022] is a novel consistency regularization method which achieved state-of-the-art performance on PLL, almost nearing supervised learning. Even so, whether IBS or ABS, both assumed that the true label is present within the candidate label set.

However, the assumption that the true labels are consistently present within the candidate label sets would be unrealistic. In existing PLL setting, the annotation for each instance is the partial labels (i.e. candidate label set) rather than the true label directly, thus significantly reduces the difficulty and cost. Against this backdrop, it poses a challenge for annotators to ensure that the true labels are present within the candidate label sets. Therefore, Unreliable Partial Label Learning (UPLL) [Lv *et al.*, 2023] is proposed in response, which is a more general problem than existing PLL. In UPLL, it is acknowledged that the true label may not be present within the candidate label set for each instance. Since that, it significantly reduces the difficulty and cost associated with data annotation. Above all, UPLL addresses the issue of labeling instances that are difficult to distinguish. Hence, UPLL could be deemed a more prevalent and valuable problem.

Despite existing PLL methods achieved promising performance, suffering from unreliable partial labeling, current PLL methods encounter numerous challenges when applied to UPLL. It will exhibit a significant decline in performance on UPLL datasets, particularly for high unreliable rates. RABS

*Corresponding authors.

[Lv *et al.*, 2023] has demonstrated that bounded loss functions have the ability to fit the ground-truth label against the interference of unreliability and other candidates. However, it would fail on high unreliability levels or high partial levels. This urges us to design an efficient method to manipulate high unreliability problem.

Motivated by this consideration, a framework named Unreliable Partial Label Learning with Recursive Separation (UPLLRS) is devised to issue this puzzle. In this paper, a novel separation method named Recursive Separation (RS) is proposed for known unreliable rate scenes to separate unreliable samples and reliable samples. However, it is limited in the real world because the real unreliable rate is difficult to know. In order to tackle this problem, more generally, we design a self-adaptive strategy for RS algorithm which could fit unknown unreliability rate. Pilot experiments have demonstrated the effectiveness of self-adaptable RS algorithm. After that, we combine a label disambiguation strategy with semi-supervised learning techniques in the second stage of UPLLRS. Experiments show that our method achieve state-of-the-art results on the UPLL datasets. Our contributions can be summarized as follows:

- A self-adaptive recursive separation algorithm is proposed for effectively separating raw dataset into a reliable subset and an unreliable subset.
- A two-stage framework is proposed for inducing the predictive model, based on the self-adaptive RS strategy. Upon obtaining both the reliable subset and unreliable subset, the disambiguation strategy utilizes the reliable subset for learning while incorporating information from the unreliable subset by the semi-supervised technique.
- The UPLLRS framework is versatile, capable of handling both image and non-image datasets. Utilizing data augmentation techniques on image datasets, the performance will be further enhanced.

The rest of this paper is organized as follows. First, we briefly review related works on partial label learning. Second, the details of the proposed UPLLRS are introduced. Third, we present the results of the comparative experiments, followed by the final conclusion.

2 Related Work

Partial label learning deals with the problem that the true label of each instance resides in the candidate label set. Many algorithms have been proposed to tackle this problem, with existing PLL methods broadly classified into classical and deep learning approaches.

In classical PLL, label disambiguation is based on averaging or identification. In averaging-based methods, the candidate label set and non-candidate label set are treated the same[Hüllermeier and Beringer, 2006; Cour *et al.*, 2011; Zhang and Yu, 2015]. For example, [Cour *et al.*, 2011] discriminated candidate labels and non-candidate labels with a convex loss. But identification-based methods progressively refine labels in the candidate set during the model training[Chen *et al.*, 2013; Yu and Zhang, 2016]. [Yu and

Zhang, 2016] optimized the constraint on the maximum margin between maximum modeling output of candidate labels and that of other labels.

However, the model output of averaging-based methods often overwhelms the true label, resulting in low accuracy. As a result, many identification-based algorithms have been devised in recent years[Feng and An, 2019; Gong *et al.*, 2017; Lyu *et al.*, 2019; Tang and Zhang, 2017; Xu *et al.*, 2019]. Nevertheless, these classical methods often have a bottleneck due to the restriction of the linear model.

Given the success of deep neural network-based methods in classification tasks, a proliferation of PLL approaches incorporating deep neural networks have emerged. [Yao *et al.*, 2020a] designed two regularization techniques in the training with ResNet representing the first exploration of deep PLL. [Yao *et al.*, 2020b], referring to the idea of co-training, trained two networks to interact with each other for label disambiguation. Concurrently, a progressive method proposed by [Lv *et al.*, 2020] progressively identified true label adopting the memorization effect of deep network. [Feng *et al.*, 2020] formalized the partial label generation process and proposed two provably consistent algorithms, risk consistent (RC) classifier and classification consistent (CC) classifier. Then, a leveraged weighted loss, which balances the contributions of candidate labels and non-candidate labels, was proposed by [Wen *et al.*, 2021]. With the development of contrastive learning, [Wang *et al.*, 2022] applied contrastive learning to PLL for effective feature representation. Recently, [Wu *et al.*, 2022] designed a consistency regularization framework in deep PLL, which gives very small performance drop compared with fully supervised learning.

However, it is common for false positive labels to be inadvertently chosen from label set, rather than being selected randomly. More specifically, it is acknowledged that each instance may not possess a uniform prior label distribution, but rather a latent label distribution which encompasses vital labeling information. Thus, [Xu *et al.*, 2021b] proposed an instance-dependent approach named VALEN, which aims to recover the latent label distribution via label enhancement [Xu *et al.*, 2023; Xu *et al.*, 2021a], leveraging it to further improve performance in real-world settings. VALEN first generates a label distribution through label enhancement, then utilizes variational inference to approximate that distribution.

In practice, despite the demonstrated empirical success of the aforementioned algorithms in the PLL task, their effectiveness is limited when the ground-truth label may not be present within the candidate label set. Therefore, the unreliable partial label learning is proposed in [Lv *et al.*, 2023], which is more general in comparison to current PLL. Furthermore, [Lv *et al.*, 2023] has proved that bounded loss, such as the Mean Absolute Error (MAE) loss and the Generalized Cross Entropy (GCE) loss, is robust against unreliability. However, the performance of RABS remains limited in the presence of high levels of unreliability.

3 Preliminaries

Let \mathcal{X} and \mathcal{Y} be feature space and label space respectively, and $p(x, y)$ be the distribution on $\mathcal{X} \times \mathcal{Y}$. Moreover, $D =$

$\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ is the training set in which \mathbf{x}_i is i -th instance and y_i is the corresponding ground-truth label, and $V = \{(\mathbf{x}_i, y_i)\}_{i=1}^k$ is the validation set including k pairs of instance \mathbf{x}_i with ground-truth label y_i .

In PLL problem, the distribution $p(\mathbf{x}, y)$ is corrupted to $p(\mathbf{x}, s)$ in which s is candidate label set satisfied $p(y_i \in s_i) = 1, \forall y_i \in \mathcal{Y}$, and the training set is corrupted to $\tilde{D} = \{(\mathbf{x}_i, s_i)\}_{i=1}^n$. The goal of PLL task is to induce a classifier from ambiguous dataset \tilde{D} .

However, in UPLL, μ called unreliable rate is the probability of true label y_i not in the candidate label set s_i , it can be expressed formally as:

$$p(y_i \in s_i) = 1 - \mu. \quad (1)$$

Such that the candidate label set s_i is corrupted to \tilde{s} that is the unreliable candidate label set. Then, the UPLL training set can be denoted as $\tilde{D} = \{(\mathbf{x}_i, \tilde{s}_i)\}_{i=1}^n$.

4 Proposed Method

In this section, we firstly propose the self-adaptive Recursive Separation (RS) algorithm which aims to differentiate reliable samples and unreliable samples effectively. After that, numerous pilot experiments were conducted which demonstrate that self-adaptive RS algorithm is effective. Finally, a framework entitled Unreliable Partial Label Learning with Recursive Separation (UPLLRS) is proposed. There are two key stages in this framework. At the beginning, the self-adaptive Recursive Separation (RS) effectively split training dataset into a reliable subset and an unreliable subset. Subsequently, in order to induce a predictive model, a disambiguation strategy is employed to progressively identify ground-truth labels while utilizing semi-supervised learning techniques in combination.

4.1 Recursive Separation

The memorization effect can be interpreted as the deep network firstly fit correct labels and then gradually fit wrong labels through the learning phase [Bai *et al.*, 2021]. In recent years, small loss trick has been demonstrated to be an effective method for addressing label noise. That inspires us to identify the reliability of samples and pay more attention to the reliable samples. Furthermore, it is discovered that the top-10% large loss samples contain more unreliable samples than any other parts after a few epochs of training, as Figure 1 shows. This motivated us to progressively take unreliable partial samples away from the training set by iteratively excluding top- γ large loss samples. Based on this idea, we introduce a multi-class classifier $f(\cdot; \theta)$ with parameters θ . More specifically, the training phase of recursive separation task is optimizing the following classical multi-class classification objective function:

$$\arg \min_{\theta} \frac{1}{n} \sum_{i=1}^n \mathcal{L}_{\text{RS}}(f(\mathbf{x}_i; \theta), s_i), \quad (2)$$

where \mathcal{L}_{RS} is the loss function for the Recursive Separation (RS) stage. According to [Liu *et al.*, 2020; Bai *et al.*, 2021], in the early-learning stage, the gradient direction of cross-entropy loss is close to the correct optimization direction. It

Algorithm 1 Self-adaptive RS Algorithm

Input: Separation network $f(\cdot; \theta)$ with trainable parameters θ ; Unreliable partial label training set $\tilde{D} = \{(\mathbf{x}_i, \tilde{s}_i)\}_{i=1}^n$ and validation set $V = \{(\mathbf{x}_i, y_i)\}_{i=1}^k$; Small epochs β for each separation step; Separation rate γ ; RS patience φ and max separation step λ .

Output: Reliable subset $\tilde{D}_R^\lambda = \{(\mathbf{x}_i, \tilde{s}_i)\}_{i=1}^m$ and unreliable subset $\tilde{D}_U^\lambda = \{(\mathbf{x}_i)\}_{i=1}^{n-m}$.

```

1: Let  $\varphi_{\text{curr}} \leftarrow 0$  and  $\text{Acc}_V \leftarrow 0$ ;
2: for  $i \leftarrow 1$  to  $\lambda$  do
3:   Randomly initialize  $\theta_0^i$ ;
4:   for  $j \leftarrow 1$  to  $\beta$  do
5:     Train  $f(\cdot; \theta_{j-1}^i)$  using dataset  $\tilde{D}_R^i$ ;
6:     Calculate loss  $l$  according Eq. 3;
7:     Update parameters from  $\theta_{j-1}^i$  to  $\theta_j^i$ ;
8:     if  $j = \beta$  then
9:       Sort  $l$  by value in descending order;
10:      Exclude top- $\gamma$  instances from  $\tilde{D}_R^i$  and add excluded instances to  $\tilde{D}_U^i$  without labels;
11:     end if
12:   end for
13:   Evaluate  $f(\cdot; \theta_j^i)$  on dataset  $V$  and calculate accuracy  $\text{Acc}_{\text{curr}}$ ;
14:   if  $\text{Acc}_{\text{curr}} < \text{Acc}_V$  then
15:      $\varphi_{\text{curr}} \leftarrow \varphi_{\text{curr}} + 1$ ;
16:     if  $\varphi_{\text{curr}} \geq \varphi$  then
17:       break;
18:     end if
19:   else
20:      $\text{Acc}_V \leftarrow \text{Acc}_{\text{curr}}, \varphi_{\text{curr}} \leftarrow 0$ ;
21:   end if
22: end for
23: return Reliable subset  $\tilde{D}_R^\lambda$  and unreliable subset  $\tilde{D}_U^\lambda$ .

```

inspires us to choose Categorical Cross Entropy (CCE) [Lv *et al.*, 2023] loss as \mathcal{L}_{RS} under the UPLL setting, such that the objective function can be rewritten as:

$$\arg \min_{\theta} \frac{1}{n} \cdot \frac{1}{|s_i|} \sum_{i=1}^n \sum_{j \in s_i} -\log p_j(f(\mathbf{x}_i; \theta)). \quad (3)$$

Following [Lv *et al.*, 2023], if the dataset has C classes, the $p_j(f(\mathbf{x}_i; \theta))$ can be specified as:

$$p_j(f(\mathbf{x}_i; \theta)) = \frac{e^{f_j(\mathbf{x}_i; \theta)}}{\sum_{k=1}^C e^{f_k(\mathbf{x}_i; \theta)}}, \quad (4)$$

where $f(\cdot)_j$ is the output for j -th class. More specifically, $p_j(f(\cdot))$ is the j -th class probability of classifier $f(\cdot)$'s output.

Since the classifier will fit more unreliable labels after the early learning stage, the classifier $f(\cdot; \theta)$ should only be trained for several epochs β . Formally speaking, let \tilde{D}_R^i be the reliable subset of i -th separation step. At the very beginning, set $\tilde{D}_R^0 = \tilde{D}$. Then let $\tilde{D}_U = \{\mathbf{x}_i\}_{i=1}^m$ denotes instances excluded from the reliable subset. θ_j^i is the parameters for i -th step separation's j -th training epoch. Note that θ_0^i

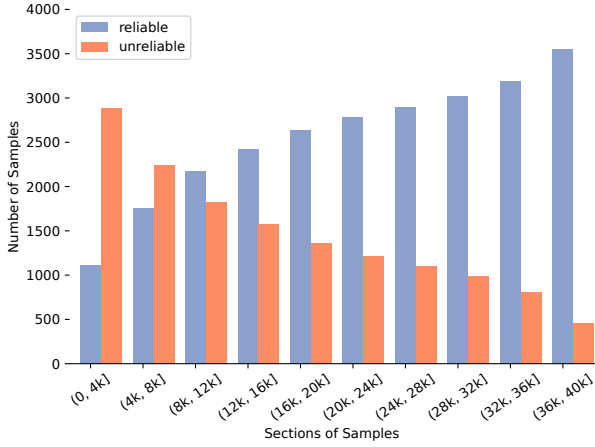


Figure 1: Number of samples for ten consecutive sections sorted by loss values in descending order. In this figure, orange bars represent the number of unreliable samples and blue bars represent the number of reliable samples. For the first section (0, 4K], which contains 4000 samples and almost 3000 samples are unreliable. But in the last section (36K, 40K], it’s the exact opposite of that.

denotes the randomly initialized parameters in i -th separation step.

The RS algorithm can be described as follows: for the i -th step, the parameters θ_0^i are randomly initialized. Then we train $f(\cdot; \theta)$ for several epochs β , get parameters θ_β^i . After that, we retrieve the final epoch (i.e. β -th epoch) training losses for each sample and sorted them by loss value in descending order. Simultaneously, instances that are correlated with the top- γ ($0 < \gamma < 1$) maximum loss values are shifted to the \tilde{D}_U^i . Following λ steps of separation, we will acquire a reliable subset \tilde{D}_R^λ and an unreliable subset \tilde{D}_U^λ .

If the unreliable rate μ on dataset is known, the λ can be estimated by μ directly. However, it is generally not feasible in real-world settings. Given this reality, a self-adaptive strategy has been devised to accommodate various levels of unreliability. Intuitively, as the count of unreliable samples descending in \tilde{D}_R^i , the accuracy on validation or test set will increase. But at the latter phase, most of the unreliable samples have been removed and samples in \tilde{D}_R^i are totally reliable. As the value of $|\tilde{D}_R^\lambda|$ goes down, the accuracy on the validation or test set diminishes. Given the need for addressing the limitations of unknown μ , we propose a self-adaptive RS algorithm that incorporates an early-stopping technique to terminate the process of separation at an appropriate time. The self-adaptive strategy dictates that, should the accuracy on the validation set cease to improve over φ consecutive epochs, the separation process will be terminated. This leads to the final determination of \tilde{D}_R^λ and \tilde{D}_U^λ , with φ representing the separation patience. The details of self-adaptive RS algorithm is exhibited in Algorithm 1.

4.2 Pilot Experiments

In order to validate the efficacy of the self-adaptive RS method put forth, a series of pilot experiments were con-




Figure 2: Real reliable rate variation on CIFAR-10 training set as separation step goes up in four different unreliable rate settings. The real reliable rate nearing 100% means the subset \tilde{D}_R^λ is almost reliable.

ducted to assess the ability of the method to identify and exclude unreliable samples, resulting in the formation of a reliable subset.

To explore this idea, we generate UPLL dataset on CIFAR-10 [Krizhevsky *et al.*, 2009] at first with four different unreliable rates 0.1, 0.2, 0.3 and 0.4. Then annotate reliable or unreliable for each sample by the true label in or not in the candidate label set respectively. The partial rate is fixed as 0.1 in the pilot experiments. MultiLayer Perceptron (MLP) is used as backbone since complex networks will overfit unreliable samples faster than plain networks. The Categorical Cross Entropy (CCE) loss [Lv *et al.*, 2023] is utilized to train the classifier. More generation process of the dataset see the section 5 below for a detailed description.

At first, we train the network for 5 epochs. In the 5-th epoch, the samples in training set are sorted by loss value in descending order. The experimental results are reported in Figure 1. The orange bars represent the number of unreliable samples and the blue bars represent the number of reliable samples. The 40K samples are divided into 10 sections and each section contains 4K samples. As is shown, the unreliable samples in the first section occupy a larger proportion than the reliable samples. But in the last section, it’s exactly the opposite of that.

It is concluded that the section with higher loss value will contain more unreliable samples than the one with lower loss value. Motivated by this finding, we try to exclude top-3% samples every 5-epoch as a separation step with four different unreliable rates $\{0.1, 0.2, 0.3, 0.4\}$. Then record the variation of real reliable rate on the training set. As shown in Figure 2, the proportion of reliable samples increases as the number of separation steps increases. That is to say, our self-adaptive RS method can effectively exclude unreliable samples and then get a highly reliable subset. With relatively low unreliability, the model is able to achieve a higher level of accuracy. In the following, a framework is designed to learn from these two subsets.

Algorithm 2 UPLLRs Algorithm with General Solution

Input: Network $g(\cdot; \omega)$ with parameters ω . Unreliable partial label dataset \tilde{D} and validation set V . Max training epochs T .
Output: Parameters ω for $g(\cdot)$.

- 1: Obtain reliable subset \tilde{D}_R^λ and unreliable partial dataset \tilde{D}_U^λ by executing Algorithm 1;
 - 2: Randomly initialize ω .
 - 3: **for** $i \leftarrow 1$ to T **do**
 - 4: Train $g(\cdot; \omega)$ from \tilde{D}_R^λ ;
 - 5: Calculate loss according Eq. 6;
 - 6: Update w_{ij} according Eq. 7;
 - 7: Use $g(\cdot; \omega)$ and get pseudo labels R which over threshold τ on dataset \tilde{D}_U^λ ;
 - 8: Add pseudo labels R and corresponding instances to reliable dataset \tilde{D}_R^λ and remove it from \tilde{D}_U^λ ;
 - 9: **end for**
 - 10: **return** ω .
-

4.3 The Overall Framework

The overall framework of UPLLRs consists of two stages. Firstly, the self-adaptive RS component recursively separates \tilde{D} , obtaining a reliable subset \tilde{D}_R^λ and an unreliable subset \tilde{D}_U^λ as a result. Then, the disambiguation strategy, in conjunction with a semi-supervised learning approach, induces the model from both \tilde{D}_R^λ and \tilde{D}_U^λ . Ultimately, we arrive at a well-trained classifier $g(\cdot; \Theta)$.

General Solution

After the recursive separation stage, the dataset is split into the reliable subset \tilde{D}_R^λ and unreliable subset \tilde{D}_U^λ . As the unreliable rate μ goes up, \tilde{D}_R will remain fewer and fewer samples while \tilde{D}_U will collect more and more unreliable samples. In order to fully leverage the information contained within the \tilde{D}_U , we employ pseudo-labeling technique for instances in \tilde{D}_U^λ . After the completion of each epoch, the model is evaluated on \tilde{D}_U^λ and high-confidence samples are added to \tilde{D}_R^λ . More specifically, the pseudo label for each instance x_i is given by:

$$u_i = \arg \max(p_g(x_i)), \quad (5)$$

where p_g denotes the model's predicted class distribution. We only retain the pseudo labels which satisfy $\max(p_g(x_i)) \geq \tau$, where τ is a threshold. It is fixed as 0.95 in our experiments.

Although the \tilde{D}_R^λ is reliable, the labels therein are ambiguous. Adopting the disambiguation method PRODEN [Lv *et al.*, 2020], the weighted loss can be written as:

$$\mathcal{L} = \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^C w_{ij} \mathcal{L}_{\text{CCE}}(g_j(x_i), s_i), \quad (6)$$

in which w_{ij} is the confidence of the j -th class being consistent with the concealed true class for the i -th instance. It is estimated by the output of classifier $g(\cdot; \omega)$, which is defined as:

$$w_{ij} = \begin{cases} g_j(x_i) / \sum_{k \in s_i} g_k(x_i) & \text{if } j \in s_i, \\ \text{otherwise,} & \end{cases} \quad (7)$$

where $g_j(\cdot)$ is the j -th coordinate of $g(\cdot)$. For initialization, the weights are uniform, i.e. if $j \in s_i$, $w_{ij} = 1/|s_i|$, otherwise $w_{ij} = 0$. The algorithm of the overall framework is illustrated in Algorithm 2.

Augmented Solution for Image Datasets

The general solution is able to handle both image and non-image datasets. As for image datasets, augmentation is an important procedure in classification tasks [Shorten and Khoshgoftaar, 2019]. Both \tilde{D}_R^λ and \tilde{D}_U^λ can employ image augmentation strategies to further enhance their performance. For the \tilde{D}_R^λ , the consistency regularization based method CR-DPLL [Wu *et al.*, 2022] is capable of being employed which takes advantage of image augmentation and achieves promising performance on PLL. To be specific, the \mathcal{L}_{PLL} can be written as:

$$\mathcal{L}_{\text{PLL}} = \mathcal{L}_{\text{Sup}}(x, s) + \pi(t) \Psi(x, s), \quad (8)$$

where $\mathcal{L}_{\text{Sup}}(x, s) = -\sum_{k \notin s} \log(1 - g_k(x))$ and $\Psi(x, s) = \sum_{z \in \mathcal{A}(x)} \text{KL}(s \| g(z))$. $\text{KL}(\cdot)$ denotes the Kullback-Leibler divergence and $\mathcal{A}(x)$ denotes the set of random augmented versions of instance x . $\pi(t) = \min\{t\pi/T', \pi\}$ is a dynamic balancing factor, where t is the current epoch and T' is a constant. More specifically, the factor is increased to π at the T' -th epoch, and thereafter maintained at a constant value of π until the end of the training. Meanwhile, the label weights are iteratively updated every epoch, more details are presented in Appendix A.1.

As for \tilde{D}_U^λ , a semi-supervised learning method [Sohn *et al.*, 2020] can be leveraged to extract potential valuable information in the unreliable instances. Specifically, for the images in the unreliable subset \tilde{D}_U^λ , the pseudo labels generated by the model's prediction where the images is weakly augmented. Next, we selectively preserve the samples whose pseudo labels satisfy the condition of $\max(p_g(x_i^w)) \geq \tau$. Then the model is trained to predict the pseudo labels when fed with a strongly-augmented version of the same image. Thus, the loss function for unreliable subset \tilde{D}_U^λ takes the following form:

$$\mathcal{L}_U = \frac{1}{n-m} \sum_{i=1}^{n-m} \mathbb{1}(\max(p_g(x_i^w)) \geq \tau) \mathcal{L}_{\text{CE}}(g(x_i^w), g(x_i^s)), \quad (9)$$

where x_i^w and x_i^s are the weak and strong augmentation of x_i respectively. $\mathbb{1}(\cdot)$ is an indicator function and \mathcal{L}_{CE} represents cross-entropy loss. Consequently, the objective of UPLLRs is as follows:

$$\mathcal{L} = \mathcal{L}_{\text{PLL}} + \xi \mathcal{L}_U, \quad (10)$$

where ξ is a scalar hyperparameter.

5 Experiments

A comprehensive set of experiments were conducted to evaluate the performance of our method under varying levels of partial and unreliable labeling. The results demonstrate that our approach achieves state-of-the-art accuracy on tasks involving UPLL.

5.1 Datasets and Implementation Details

Datasets

We utilize two commonly employed image datasets, CIFAR-10 and CIFAR-100 [Krizhevsky *et al.*, 2009], as the basis for

Dataset	η	μ	Ours	RABS	PiCO	CR-DPLL	PRODEN	RC	CC	LWS
CIFAR-10	0.1	0.1	95.16 \pm 0.10%	83.87 \pm 0.29%	91.35 \pm 0.14%	93.49 \pm 0.26%	79.77 \pm 0.61%	79.96 \pm 0.46%	78.91 \pm 0.61%	85.83 \pm 0.76%
	0.1	0.3	94.65 \pm 0.23%	77.75 \pm 0.62%	87.66 \pm 0.22%	90.65 \pm 0.20%	67.80 \pm 1.38%	69.46 \pm 1.02%	67.52 \pm 2.11%	19.95 \pm 4.22%
	0.1	0.5	93.12 \pm 0.92%	65.09 \pm 0.57%	82.47 \pm 0.38%	85.65 \pm 0.38%	51.07 \pm 1.49%	54.75 \pm 1.57%	52.37 \pm 2.95%	16.65 \pm 1.35%
	0.3	0.1	94.32 \pm 0.21%	53.13 \pm 0.90%	90.50 \pm 0.24%	92.92 \pm 0.15%	77.12 \pm 0.32%	75.39 \pm 0.31%	75.37 \pm 0.61%	83.92 \pm 0.35%
	0.3	0.3	93.85 \pm 0.31%	41.61 \pm 2.11%	86.37 \pm 0.37%	88.80 \pm 0.19%	62.06 \pm 0.69%	61.87 \pm 1.63%	62.91 \pm 1.20%	78.33 \pm 0.68%
	0.3	0.5	91.16 \pm 0.67%	30.33 \pm 1.61%	79.87 \pm 0.51%	82.06 \pm 0.34%	44.38 \pm 0.97%	47.13 \pm 0.62%	45.75 \pm 2.31%	24.16 \pm 2.34%
	0.5	0.1	92.47 \pm 0.19%	31.62 \pm 2.31%	89.48 \pm 0.38%	91.88 \pm 0.32%	73.30 \pm 0.07%	68.17 \pm 0.55%	71.03 \pm 0.33%	70.46 \pm 3.00%
	0.5	0.3	91.55 \pm 0.38%	27.88 \pm 2.58%	84.48 \pm 0.33%	86.78 \pm 0.54%	57.25 \pm 0.98%	54.55 \pm 0.64%	54.69 \pm 1.64%	58.31 \pm 4.76%
	0.5	0.5	89.56 \pm 0.50%	24.48 \pm 2.77%	74.68 \pm 1.21%	78.31 \pm 0.41%	42.99 \pm 0.80%	42.43 \pm 1.17%	36.96 \pm 1.78%	40.23 \pm 4.16%
CIFAR-100	0.01	0.1	75.73 \pm 0.41%	27.38 \pm 1.42%	67.94 \pm 0.52%	74.22 \pm 0.41%	55.68 \pm 0.49%	56.08 \pm 0.37%	55.35 \pm 0.76%	5.37 \pm 0.61%
	0.01	0.3	71.72 \pm 0.39%	17.56 \pm 0.73%	62.12 \pm 0.35%	68.56 \pm 0.37%	45.31 \pm 0.63%	44.80 \pm 1.20%	44.95 \pm 0.74%	3.61 \pm 0.91%
	0.01	0.5	66.40 \pm 0.21%	11.73 \pm 0.62%	54.84 \pm 0.40%	61.93 \pm 0.38%	32.87 \pm 0.90%	32.55 \pm 1.17%	33.62 \pm 0.81%	3.17 \pm 0.72%
	0.05	0.1	74.73 \pm 0.24%	31.65 \pm 0.79%	66.67 \pm 0.46%	73.34 \pm 0.43%	52.05 \pm 0.90%	50.04 \pm 0.27%	52.02 \pm 0.35%	16.11 \pm 4.07%
	0.05	0.3	70.31 \pm 0.22%	21.45 \pm 0.74%	59.01 \pm 0.61%	66.79 \pm 0.75%	37.81 \pm 0.89%	25.06 \pm 0.75%	40.24 \pm 0.84%	8.49 \pm 0.92%
	0.05	0.5	64.78 \pm 0.53%	15.08 \pm 1.18%	46.81 \pm 0.69%	59.09 \pm 0.76%	20.84 \pm 1.25%	19.93 \pm 0.92%	26.08 \pm 0.66%	6.65 \pm 0.66%
	0.1	0.1	73.20 \pm 0.50%	25.55 \pm 1.55%	45.44 \pm 1.68%	72.08 \pm 0.52%	44.07 \pm 0.47%	38.70 \pm 1.52%	47.81 \pm 0.90%	49.91 \pm 0.97%
	0.1	0.3	68.60 \pm 0.25%	16.99 \pm 2.06%	35.89 \pm 1.48%	64.70 \pm 0.45%	25.66 \pm 0.58%	21.26 \pm 0.63%	34.02 \pm 0.76%	18.11 \pm 2.83%
	0.1	0.5	60.66 \pm 0.75%	10.80 \pm 0.62%	22.57 \pm 1.07%	52.34 \pm 0.62%	13.61 \pm 0.63%	12.89 \pm 0.62%	20.63 \pm 0.63%	9.52 \pm 0.46%

Table 1: Test accuracy (mean \pm std) on CIFAR-10 and CIFAR-100 synthesized dataset. The best results are highlighted in bold.

	Dermatology					20Newsgroups				
	η	0.1	0.1	0.3	0.3	0.1	0.1	0.3	0.3	
	μ	0.3	0.5	0.3	0.5	0.3	0.5	0.3	0.5	
Ours		96.06 \pm 1.31%	89.75 \pm 1.31%	91.80 \pm 2.74%	87.87 \pm 6.52%	72.27 \pm 1.55%	65.41 \pm 0.96%	61.47 \pm 1.13%	47.98 \pm 1.51%	
RABS		78.36 \pm 6.33%	47.86 \pm 6.67%	58.69 \pm 5.42%	46.88 \pm 8.95%	64.18 \pm 1.00%	50.99 \pm 0.79%	31.97 \pm 1.09%	23.45 \pm 0.53%	
PRODEN		82.95 \pm 4.34%	62.29 \pm 10.52%	77.70 \pm 3.96%	60.98 \pm 8.88%	64.08 \pm 0.43%	50.69 \pm 1.36%	58.79 \pm 1.03%	42.96 \pm 0.90%	
RC		79.34 \pm 4.82%	61.63 \pm 7.30%	76.39 \pm 7.93%	54.75 \pm 7.86%	63.23 \pm 0.70%	48.33 \pm 0.84%	56.09 \pm 0.71%	39.43 \pm 1.06%	
CC		83.93 \pm 3.34%	60.65 \pm 9.72%	81.97 \pm 5.18%	55.41 \pm 8.52%	62.39 \pm 1.05%	48.10 \pm 0.39%	54.55 \pm 0.88%	37.19 \pm 1.38%	
LWS		83.28 \pm 4.90%	74.42 \pm 12.33%	77.05 \pm 3.28%	63.93 \pm 7.33%	40.17 \pm 4.64%	24.99 \pm 2.16%	11.20 \pm 1.08%	9.14 \pm 0.64%	

Table 2: Test accuracy (mean \pm std) on UCI synthesized dataset.

synthesizing our UPLL dataset. Besides, we also utilize two additional datasets Dermatology and 20Newsgroups from UCI machine learning Repository [Dua and Graff, 2017] to further validate the effectiveness of our proposed method. In our experiments, the datasets are partitioned into training, validation, test set in a 4:1:1 ratio. Further elaboration can be found in Appendix A.2.

Following the confusing strategy in [Lv *et al.*, 2023], the ground-truth labels in the raw dataset are corrupted initially and then generate partial labels by the flipping process. That is to say, for an instance x with the ground-truth label $y = i, i \in \mathcal{Y}$, it has a fixed probability $1 - \mu$ do not make any operation. But it has a probability κ to flip into j , where $j \in \mathcal{Y}, j \neq i, \kappa = \mu / (C - 1)$. The unreliable label is called \tilde{y}_i . Subsequently, \tilde{y}_i is considered as the true label for generating the candidate label set, employing a uniform partial labeling with probability η in accordance with the approach presented in [Lv *et al.*, 2020], where η denotes the partial rate.

5.2 Baselines

In order to demonstrate the efficacy of our proposed method and to gain insight into its underlying characteristics, we conduct comparisons with seven benchmark methods including one UPLL method and six state-of-the-art PLL methods: 1) RABS [Lv *et al.*, 2023]: An unreliable PLL method that proved the robustness of Average-Based Strategy (ABS) with bounded loss function in mitigating the impact of unreliability. In our experiment, the Mean Average Error (MAE) loss is

chosen as the baseline. 2) PiCO [Wang *et al.*, 2022]: A PLL method combines the idea of contrastive learning and class prototype-based label disambiguation method. 3) CR-DPLL [Wu *et al.*, 2022]: A deep PLL method based on consistency regularization. 4) PRODEN [Lv *et al.*, 2020]: A PLL method which progressively identifies true labels in candidate label sets. 5) RC [Feng *et al.*, 2020]: A risk-consistent method for PLL which employs importance re-weighting strategy. 6) CC [Feng *et al.*, 2020]: A classifier-consistent method for PLL using transition matrix to form an empirical risk estimator. 7) LWS [Wen *et al.*, 2021]: A PLL method utilizing Leveraged weighted (LW) loss which balances the trade-off between losses on partial labels and others.

Note that two partial label learning methods PiCO [Wang *et al.*, 2022] and CR-DPLL [Wu *et al.*, 2022] are not suitable on the Dermatology and 20Newsgroups. More details can be found in Appendix A.2.

Implementation Details

For the first stage (i.e. self-adaptive RS), a 5-layer perceptron (MLP) is utilized to separate samples with CCE [Lv *et al.*, 2023] loss. The learning rate is 0.1, 0.18, 0.1; small epochs $\beta = 5, 6, 5$; separation rate $\gamma = 0.03, 0.005, 0.03$; on the CIFAR-10, CIFAR-100 and UCI datasets respectively. Max separation step $\lambda = \lfloor \log_{1-\gamma} 0.3 \rfloor$. As for the second stage, we employ different backbones for different datasets. On CIFAR-10 dataset and CIFAR-100 dataset, we use WideResNet28 \times 2 [Zagoruyko and Komodakis, 2016] as the predictive model, and we employ the Augmented So-

Ablation	RS	\tilde{D}_U^λ	CIFAR-10 ($\eta = 0.3$)			CIFAR-100 ($\eta = 0.05$)		
			$\mu = 0.1$	$\mu = 0.3$	$\mu = 0.5$	$\mu = 0.1$	$\mu = 0.3$	$\mu = 0.5$
UPLLRs	✓	✓	94.32 \pm 0.21%	93.85 \pm 0.31%	91.16 \pm 0.67%	74.73 \pm 0.24%	70.31 \pm 0.22%	64.78 \pm 0.53%
UPLLRs w/o \tilde{D}_U^λ	✓	✗	93.07 \pm 0.08%	92.48 \pm 0.27%	89.81 \pm 0.39%	74.35 \pm 0.52%	70.38 \pm 0.50%	64.56 \pm 0.40%
UPLLRs w/o RS	✗	✗	92.92 \pm 0.15%	88.80 \pm 0.19%	82.06 \pm 0.34%	73.34 \pm 0.43%	66.79 \pm 0.75%	59.09 \pm 0.76%

 Table 3: The impact of RS and Unreliable Subset \tilde{D}_U^λ on accuracy (mean \pm std).

lution. On the UCI datasets, a 5-layer perceptron (MLP) is employed, and we utilize the General Solution. The learning rate is $5e - 2$ and the weight decay is $1e - 3$; ξ is set as 2 on CIFAR-10 and 0.3 on CIFAR-100. We implement the data augmentation technique following the "strong augmentation" in CR-DPLL [Wu *et al.*, 2022]. This processing is applied on Ours, PRODEN, RC, CC, LWS. As for PiCO and CR-DPLL, the augmentation setups followed recommended setting in the previous works.

The optimizer in our experiment is Stochastic Gradient Descent (SGD) [Robbins and Monro, 1951] in which momentum is set as 0.9. For the learning rate scheduler, we use a cosine learning rate decay [Loshchilov and Hutter, 2016]. Otherwise, each model is trained with maximum epochs $T = 500$ and employs early stopping strategy with patience 25. In other words, if the accuracy does not rise in validation set V for 25 epochs, the training process will be stopped. All experiments are conducted on NVIDIA RTX 3090. What's more, the implementation of our method is based on PyTorch [Paszke *et al.*, 2019] framework. We report final performance using the test accuracy corresponding to the best accuracy on validation set for each run. Finally, we report the mean and standard deviation based on five independent runs with different random seeds.

5.3 Experiment Results

Table 1 reports the experimental results on CIFAR-10 and CIFAR-100 synthesized datasets. As is shown, our UPLLRs method outperforms all compared methods. The improvements are particularly pronounced in scenarios with high levels of unreliability. Take $\eta = \{0.1, 0.3, 0.5\}$, $\mu = 0.5$ on CIFAR-10 dataset as an instance, our method improves by 7.47%, 9.1%, 11.25% respectively compared with the second-best methods. It is worth noting that our method exhibits a minimal decline in accuracy as the unreliable rate μ increases. For example, for $\eta = 0.5$, the accuracy for $\mu = 0.1$ is 92.47% and $\mu = 0.5$ is 89.56%, only 2.91% accuracy drop. In contrast, the second-best method's accuracy drop up to 13.57%. UPLLRs also achieves the best performance and significantly outperforms other compared methods on the CIFAR-100 synthesized dataset. For the settings with $\eta = 0.1$, $\mu = 0.5$, our UPLLRs also achieves 60.66% outperforming second-best 8.32%. Conversely, other methods either exhibit poor performance or fail to converge.

We further evaluate the performance of UPLLRs on non-image datasets Dermatology and 20Newsgroups. Table 2 reports the experimental results on it. Our method demonstrates a clear advantage and surpasses all the methods evaluated in comparison. For instance, in the condition of $\eta = 0.3$, $\mu = \{0.3, 0.5\}$ on Dermatology, our method exhibited a 9.83%,

23.94% over the second-best method respectively. Furthermore, during experimentation with $\eta = 0.1$, our method exhibited a drop of 11.29% when varying the μ from 0.3 to 0.5. Hence, our method has noticeable resistance to unreliability. In contrast, other methods exhibit a significant decline.

5.4 Ablation Study

In this subsection, we present the results of our ablation study which serve to demonstrate the efficacy of the components of our UPLLRs method: RS and Unreliable Subset \tilde{D}_U^λ . The experiments are conducted on CIFAR-10 dataset with $\eta = 0.3$, $\mu = \{0.1, 0.3, 0.5\}$ and CIFAR-100 dataset with $\eta = 0.05$, $\mu = \{0.1, 0.3, 0.5\}$. Other hyperparameter settings are consistent with those utilized in the primary experiments. Besides, analysis on hyperparameter ξ and γ are elaborated in detail in the Appendix A.3.

Here, we conduct ablation studies on the individual components to investigate their contributions. Two variants are selected: 1) Without RS. That implies that the corrupted dataset will not be partitioned into subsets, but rather utilized directly to induce the final classifier. 2) Without unreliable subset \tilde{D}_U^λ . It can be stated that the final classifier is directly trained on the reliable subset. All other parameters are held constant as in the primary experiment. As shown in Table 3, It is apparent that the contribution of RS surpasses that of \tilde{D}_U^λ . Take CIFAR-10 with $\mu = 0.5$ as an instance, it was observed that the variant without the utilization of the \tilde{D}_U^λ experienced a mere 1.35% decline in performance compared to the full UPLLRs model. However, when comparing the variant without the self-adaptive RS to the variant without the \tilde{D}_U^λ , a significant decline of 7.75% was observed. Furthermore, while utilizing the \tilde{D}_U^λ on the CIFAR-100 dataset resulted in a slight increase, this can likely be attributed to the lower accuracy of pseudo-label generation from \tilde{D}_U^λ , as CIFAR-100 has a substantially larger number of classes than CIFAR-10.

6 Conclusion

In this work, we propose a novel two-stage framework named Unreliable Partial Label Learning with Recursive Separation (UPLLRs). First, the self-adaptive recursive separation strategy is proposed to separate the training set into a reliable subset and an unreliable subset. Second, a disambiguation strategy progressively identifies ground-truth labels in the reliable subset. Meanwhile, the semi-supervised learning techniques are employed for the unreliable subset. Experimental results demonstrate that our method attains state-of-the-art performance, particularly exhibiting robustness in scenarios with high levels of unreliability.

Acknowledgments

This work is supported by National Key R&D Program of China (2018AAA0100104), the National Science Foundation of China (62206050, 62125602, and 62076063), China Postdoctoral Science Foundation (2021M700023), Jiangsu Province Science Foundation for Youths (BK20210220), Young Elite Scientists Sponsorship Program of Jiangsu Association for Science and Technology (TJ-2022-078), and the Big Data Computing Center of Southeast University.

References

- [Bai *et al.*, 2021] Yingbin Bai, Erkun Yang, Bo Han, Yanhua Yang, Jiatong Li, Yinian Mao, Gang Niu, and Tongliang Liu. Understanding and improving early stopping for learning with noisy labels. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Chen *et al.*, 2013] Yi-Chen Chen, Vishal M Patel, Jaisankar K Pillai, Rama Chellappa, and P Jonathon Phillips. Dictionary learning from ambiguously labeled data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 353–360, 2013.
- [Chen *et al.*, 2018] Ching-Hui Chen, Vishal M Patel, and Rama Chellappa. Learning from ambiguously labeled face images. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 40(07):1653–1667, 2018.
- [Cour *et al.*, 2011] Timothee Cour, Ben Sapp, and Ben Taskar. Learning from partial labels. *The Journal of Machine Learning Research*, 12:1501–1536, 2011.
- [Dua and Graff, 2017] Dheeru Dua and Casey Graff. UCI machine learning repository. Available at <http://archive.ics.uci.edu/ml>, 2017.
- [Feng and An, 2019] Lei Feng and Bo An. Partial label learning with self-guided retraining. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 3542–3549, 2019.
- [Feng *et al.*, 2020] Lei Feng, Jiaqi Lv, Bo Han, Miao Xu, Gang Niu, Xin Geng, Bo An, and Masashi Sugiyama. Provably consistent partial-label learning. *Advances in Neural Information Processing Systems*, 33:10948–10960, 2020.
- [Gong *et al.*, 2017] Chen Gong, Tongliang Liu, Yuanyan Tang, Jian Yang, Jie Yang, and Dacheng Tao. A regularization approach for instance-based superset label learning. *IEEE transactions on cybernetics*, 48(3):967–978, 2017.
- [Hüllermeier and Beringer, 2006] Eyke Hüllermeier and Jürgen Beringer. Learning from ambiguously labeled examples. *Intelligent Data Analysis*, 10(5):419–439, 2006.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Master’s thesis, Dept. of Comp. Sci., University of Toronto*, 2009.
- [Liu and Dietterich, 2012] Liping Liu and Thomas Dietterich. A conditional multinomial mixture model for superset label learning. *Advances in neural information processing systems*, 25, 2012.
- [Liu *et al.*, 2020] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. *Advances in neural information processing systems*, 33:20331–20342, 2020.
- [Loshchilov and Hutter, 2016] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [Luo and Orabona, 2010] Jie Luo and Francesco Orabona. Learning from candidate labeling sets. *Advances in neural information processing systems*, 23, 2010.
- [Lv *et al.*, 2020] Jiaqi Lv, Miao Xu, Lei Feng, Gang Niu, Xin Geng, and Masashi Sugiyama. Progressive identification of true labels for partial-label learning. In *International Conference on Machine Learning*, pages 6500–6510. PMLR, 2020.
- [Lv *et al.*, 2023] Jiaqi Lv, Biao Liu, Lei Feng, Ning Xu, Miao Xu, Bo An, Gang Niu, and Xin Geng. On the robustness of average losses for partial-label learning. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, in press, 2023.
- [Lyu *et al.*, 2019] Gengyu Lyu, Songhe Feng, Tao Wang, Congyan Lang, and Yidong Li. Gm-pll: graph matching based partial label learning. *IEEE Transactions on Knowledge and Data Engineering*, 33(2):521–535, 2019.
- [Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [Robbins and Monro, 1951] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [Shorten and Khoshgoftaar, 2019] Connor Shorten and Taghi M. Khoshgoftaar. A survey on image data augmentation for deep learning. *J. Big Data*, 6:60, 2019.
- [Sohn *et al.*, 2020] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020.
- [Tang and Zhang, 2017] Cai-Zhi Tang and Min-Ling Zhang. Confidence-rated discriminative partial label learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- [Wang *et al.*, 2022] Haobo Wang, Ruixuan Xiao, Yixuan Li, Lei Feng, Gang Niu, Gang Chen, and Junbo Zhao. Pico: Contrastive label disambiguation for partial label learning. *arXiv preprint arXiv:2201.08984*, 2022.

- [Wen *et al.*, 2021] Hongwei Wen, Jingyi Cui, Hanyuan Hang, Jiabin Liu, Yisen Wang, and Zhouchen Lin. Leveraged weighted loss for partial label learning. In *International Conference on Machine Learning*, pages 11091–11100. PMLR, 2021.
- [Wu *et al.*, 2022] Dong-Dong Wu, Deng-Bao Wang, and Min-Ling Zhang. Revisiting consistency regularization for deep partial label learning. In *International Conference on Machine Learning*, pages 24212–24225. PMLR, 2022.
- [Xu *et al.*, 2019] Ning Xu, Jiaqi Lv, and Xin Geng. Partial label learning via label enhancement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5557–5564, 2019.
- [Xu *et al.*, 2021a] Ning Xu, Yun-Peng Liu, and Xin Geng. Label enhancement for label distribution learning. *IEEE Transactions on Knowledge and Data Engineering*, 33(4):1632–1643, 2021.
- [Xu *et al.*, 2021b] Ning Xu, Congyu Qiao, Xin Geng, and Min-Ling Zhang. Instance-dependent partial label learning. *Advances in Neural Information Processing Systems*, 34:27119–27130, 2021.
- [Xu *et al.*, 2023] Ning Xu, Jun Shu, Renyi Zheng, Xin Geng, Deyu Meng, and Min-Ling Zhang. Variational label enhancement. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (01):1–15, 2023.
- [Yao *et al.*, 2020a] Yao Yao, Jiehui Deng, Xiuhua Chen, Chen Gong, Jianxin Wu, and Jian Yang. Deep discriminative cnn with temporal ensembling for ambiguously-labeled image classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12669–12676, 2020.
- [Yao *et al.*, 2020b] Yao Yao, Chen Gong, Jiehui Deng, and Jian Yang. Network cooperation with progressive disambiguation for partial label learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 471–488. Springer, 2020.
- [Yu and Zhang, 2016] Fei Yu and Min-Ling Zhang. Maximum margin partial label learning. In *Asian conference on machine learning*, pages 96–111. PMLR, 2016.
- [Zagoruyko and Komodakis, 2016] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016.
- [Zeng *et al.*, 2013] Zinan Zeng, Shijie Xiao, Kui Jia, Tsung-Han Chan, Shenghua Gao, Dong Xu, and Yi Ma. Learning by associating ambiguously labeled images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 708–715, 2013.
- [Zhang and Yu, 2015] Min-Ling Zhang and Fei Yu. Solving the partial label learning problem: An instance-based approach. In *Twenty-fourth international joint conference on artificial intelligence*, 2015.