

Towards Effective Visual Representations for Partial-Label Learning

Shiyu Xia^{1,*} Jiaqi Lv^{2,*} Ning Xu¹ Gang Niu^{2,1} Xin Geng^{1,†}

¹Southeast University ²RIKEN Center for Advanced Intelligence Project

{shiyu_xia, xning, xgeng}@seu.edu.cn, {is.jiaqi.lv, gang.niu.ml}@gmail.com

Abstract

Under partial-label learning (PLL) where, for each training instance, only a set of ambiguous candidate labels containing the unknown true label is accessible, contrastive learning has recently boosted the performance of PLL on vision tasks, attributed to representations learned by contrasting the same/different classes of entities. Without access to true labels, positive points are predicted using pseudo-labels that are inherently noisy, and negative points often require large batches or momentum encoders, resulting in unreliable similarity information and a high computational overhead. In this paper, we rethink a state-of-the-art contrastive PLL method PiCO [24], inspiring the design of a simple framework termed PaPi (**P**artial-label learning with a guided **P**rotypical classifier), which demonstrates significant scope for improvement in representation learning, thus contributing to label disambiguation. PaPi guides the optimization of a prototypical classifier by a linear classifier with which they share the same feature encoder, thus explicitly encouraging the representation to reflect visual similarity between categories. It is also technically appealing, as PaPi requires only a few components in PiCO with the opposite direction of guidance, and directly eliminates the contrastive learning module that would introduce noise and consume computational resources. We empirically demonstrate that PaPi significantly outperforms other PLL methods on various image classification tasks.

1. Introduction

The excellent performance of modern deep neural networks (DNNs) is attributed to the large-scale fully supervised training data, but the requirement for high-quality data poses a challenge for the practical application. As a result, non-expert but cheap labelers are often resorted as an appealing substitute, which inevitably leads to low-quality labeled data due to their expertise limitation. A typical sit-




Figure 1. An image of “Eagle” in PLL is equipped with a candidate label set. PLL learns from such ambiguous supervision, in contrast to its supervised counterpart where only the true label is chosen.

uation is that the labelers have difficulty in making an accurate judgement about an instance from multiple ambiguous labels, and therefore choose multiple likely ones. For example, in Fig. 1, it can be difficult for labelers without specialist knowledge to identify an Eagle from a Hawk, so both “Eagle” and “Hawk” are labeled as possible candidates for the true label. Learning from such training instances with a set of possible candidate labels, where only one fixed but unknown is true, is known as *partial-label learning* (PLL) [3, 4, 15, 22–24, 32, 38, 40]. This problem naturally arises in various important applications in the real world such as web mining [13] and image annotation [1].

Research into PLL dates back some twenty years and a number of practical approaches have been proposed, which can be divided into identification-based strategies [9, 20, 32, 38, 41] and average-based strategies [3, 8], depending on how they treat candidate labels. Recently, DNNs bring the research of PLL into a new era [4, 15, 26–28, 36, 37], among which PiCO [24] has achieved state-of-the-art performance on multiple benchmarks. It introduces a *contrastive learning module* into PLL that uses predictions of one *linear classifier* to select pseudo positive samples for each anchor point and maintains a queue of negative samples. Meanwhile, a momentum encoder is used to improve consistency. In addition, PiCO adds a *prototypical classifier module* (called prototype-based in the original) to guide the update of the linear classifier, which is based on the idea that there exists an embedding space where points from the same class cluster around its prototype [19]. PiCO claims credit for its success to the mutual benefit of contrastive learning and prototype-based label disambiguation.

* Equal contributions. † Corresponding author.

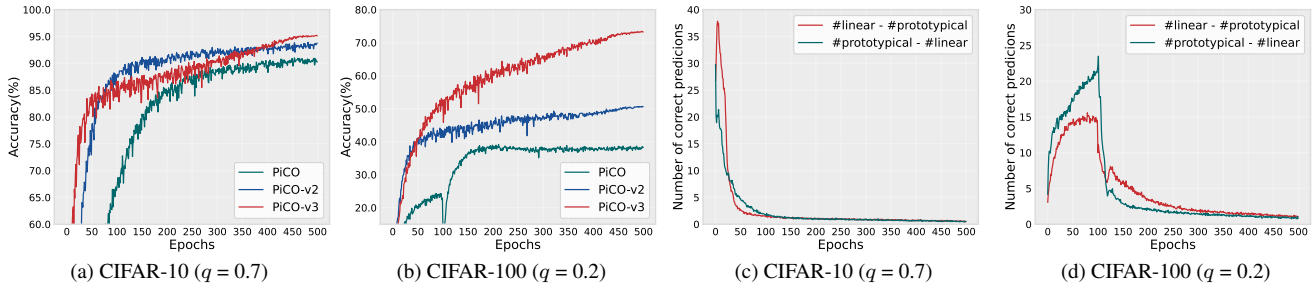


Figure 2. Visualizations of impacts of the unreliability of pseudo positives and the improper direction of disambiguation guidance in PiCO. In (a)-(b), PiCO-v2 means positives are selected based on fully supervised information, *i.e.*, true labels are known by the contrastive learning module. Further, PiCO-v3 removes the guidance of prototypical classifier to linear classifier, such that the linear classifier performs self-teaching. The red lines in (c)-(d) indicate the number of samples that were correctly classified by linear classifier and incorrectly classified by prototypical classifier per mini-batch, and the green lines are the opposite. The first 100 epochs shown in (d) are in a warm-up period. q means the flipping probability of each incorrect label, which will be introduced in Sec. 3.1.

In this paper, we rethink the two modules in PiCO and empirically point out that they do not work as well in practice as one might think due to the unreliability of the pseudo positives and the improper direction of disambiguation guidance. Fig. 2a and Fig. 2b show accuracy of three versions of PiCO. Fig. 2c and Fig. 2d show the performance differences between the linear and prototypical classifier during training. Fig. 2 delivers two important messages: (1) noisy pseudo-labels do lead to significant performance degradation, and (2) the phenomenon “poor teacher teaches good student” possibly happens. Specifically, the good student, the linear classifier, always made more correct predictions than its teacher, the prototypical classifier, at the beginning. In some cases, due to the forced direction of guidance, the teacher performed better than the student for a while, but soon the teacher had nothing new to teach the student, shown in Fig. 2c. And sometimes the student’s advantage was even maintained until convergence as shown in Fig. 2d. These also explain the significant improvement compared to PiCO-v2 after PiCO-v3 made the clever student perform self-teaching.

Inspired by the above observations, we propose a simple PLL framework termed **PaPi**, *i.e.*, *Partial-label learning with a guided Prototypical classifier*. PaPi directly eliminates the contrastive learning module which introduces noisy positives, and adopts the opposite direction of disambiguation guidance compared to PiCO. Specifically, PaPi produces a similarity distribution over classes for each sample based on a softmax function over distances to class-specific prototypes in a projected low-dimensional space. Afterwards, PaPi aligns the distribution with the disambiguated probability post-processed from one linear classifier prediction. Meanwhile, the linear classifier performs self-teaching wherein each stage of learning is guided by the current and previous stages. We conduct extensive experiments on multiple image classification tasks. PaPi sur-

passes the state-of-the-art PLL methods by a large margin, with a 4.57% improvement on CIFAR-100 in very difficult scenarios. Moreover, PaPi learns effective representations efficiently without using neither large batches nor a momentum encoder, where training instances from the same class are grouped into tighter clusters. Our main **contributions** are summarized as follows:

- We propose a simple PLL framework termed PaPi which explicitly encourages the representation to reflect visual similarity between categories, such that PaPi is remarkable for improving the class-level discrimination of learned representation.
- Extensive experiments on various image classification datasets with different generation processes of candidate labels demonstrate PaPi significantly outperforms state-of-the-art PLL methods.

2. Related Work

2.1. Partial-Label Learning

A plethora of PLL works purify the candidate labels with a common goal to select the most likely true label in the training phase, which are named as *identification-based* methods [16, 20, 41]. For this purpose, a maximum likelihood model is proposed in [9] using the expectation-maximization algorithm, which pioneers this main research route. In addition, some works [32, 38] treat the true label as a latent variable and identify it by leveraging the topological information from feature space. In contrast to this research route, *average-based* methods [3, 8, 40] treat all the candidate labels equally and the prediction is made by averaging their modeling outputs. However, most of the above methods need elaborately designed learning objectives and specific solutions, which make them inefficient for scaling up to large scale datasets. With the advances

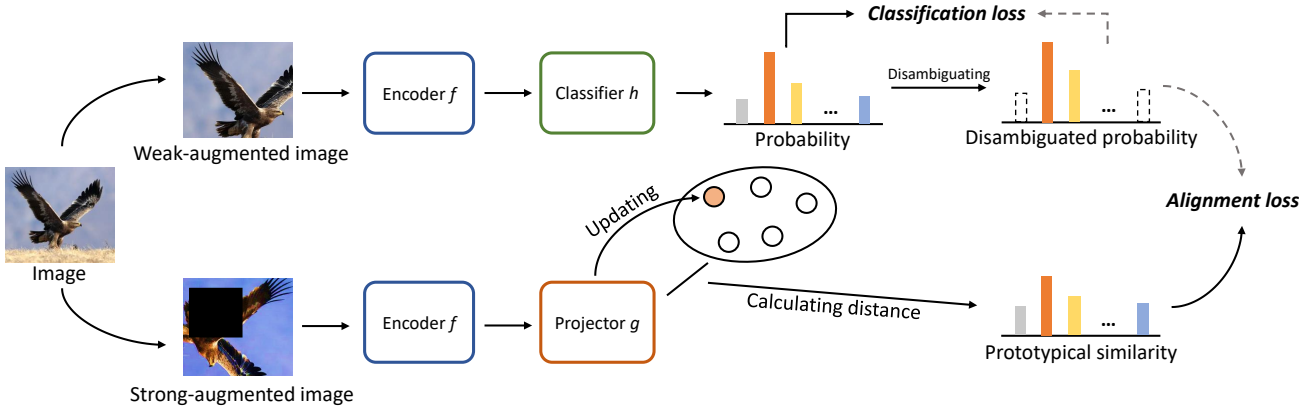


Figure 3. Illustration of PaPi. With a prototypical alignment loss term, we reduce the diversity between a prototypical similarity distribution over classes for each sample based on a softmax function over distances to class-specific prototypes in a projected low-dimensional space and the disambiguated probability post-processed from one linear classifier. Using a cross-entropy loss, we learn the linear classifier in a self-teaching fashion. Note that the dotted gray arrow indicates the stop gradient operation.

in the deep learning era, PLL has been broadly explored in recent works [4, 15, 24, 27–29, 36, 37]. PRODEN [15] accomplishes the update of the model and identification of true labels seamlessly along with a classifier-consistent risk estimator. Two methods which are risk-consistent and classifier-consistent are derived based on the uniform partial labels generation process [4]. VALEN [33] applies label enhancement [31, 34] to iteratively recover label distribution for each instance, where instance-dependent PLL is firstly explored. [14] proposes a general scenario called unreliable PLL where the ground-truth label can be unnecessarily included in the candidate label set. Consistency regularization is revisited in deep PLL [27], which employs regularization on candidate labels based on specific data augmentation techniques. In this paper, we rethink a state-of-the-art method PiCO [24] and propose a simple framework named PaPi, which demonstrates significant scope for improvement in representation learning, thus contributing to label disambiguation. Moreover, current methods cannot effectively handle both uniform and instance-dependent PLL setting, especially when they are confronted with *high ambiguity levels* and *instance-dependent ambiguity*.

2.2. Prototype Learning

Prototype learning (PL) learns a representation space where instances are imposed to be closer to its corresponding class prototype. PL has demonstrated robust performance in few-shot learning [18, 19], zero-shot learning [35], noisy-label learning [11], semi-supervised learning [6], class-imbalanced learning [25], contrastive learning [12], *etc.* For example, MoPro [11] aims to learn an embedding space where instances from the same class gather around its class prototype, which is achieved with a prototypical contrastive loss and an instance contrastive loss.

3. The Proposed Framework

In this section, we firstly introduce the preliminaries in Sec. 3.1. Then we present our PaPi in detail in Sec. 3.2. Fig. 3 gives a brief illustration of PaPi.

3.1. Preliminaries

Notations. Let \mathcal{X} be the input space, $\mathcal{Y} = \{1, 2, \dots, K\}$ be the label space with K class labels. One of the key differences between PLL and supervised learning is that the latent ground-truth label $y_i \in \mathcal{Y}$ of an instance \mathbf{x}_i is not provided but always included in a candidate label set $Y_i \subseteq \mathcal{Y}$. Following [24], we generate uniform partial labels by flipping each incorrect label $\bar{y}_i \neq y_i$ with a probability $q = P(\bar{y}_i \in Y_i | \bar{y}_i \neq y_i)$. For generating instance-dependent partial labels, we follow the same procedure in [33]. We set the flipping probability of each incorrect label j as $q_j(\mathbf{x}_i) = \frac{\hat{g}_j(\mathbf{x}_i)}{\max_{k \in \bar{Y}_i} \hat{g}_k(\mathbf{x}_i)}$ with a pre-trained DNN \hat{g} , where \bar{Y}_i is the incorrect label set of \mathbf{x}_i . Given a PLL training dataset $\mathcal{D} = \{(\mathbf{x}_i, Y_i) | 1 \leq i \leq N\}$, the goal of PLL is to learn a mapping function which can predict the true label related to the unseen inputs.

Overview. For each sample (\mathbf{x}_i, Y_i) , we generate one weakly-augmented view $(aug_1(\mathbf{x}_i), Y_i)$ and one strongly-augmented view $(aug_2(\mathbf{x}_i), Y_i)$, where $aug_1(\cdot)$ and $aug_2(\cdot)$ represent weak and strong augmentation function respectively. Then both augmented samples are separately fed into the weight-shared encoder network $f(\cdot)$, yielding a pair of representations $\mathbf{v}_i^1 = f(aug_1(\mathbf{x}_i))$ and $\mathbf{v}_i^2 = f(aug_2(\mathbf{x}_i))$. In order to improve the representation quality [2, 5], we further map \mathbf{v}_i^1 to $\mathbf{z}_i^1 = g(\mathbf{v}_i^1) \in \mathbb{R}^{d_p}$ by utilizing a projector network $g(\cdot)$, which is also operated on \mathbf{v}_i^2 to obtain \mathbf{z}_i^2 . Therefore, for each instance \mathbf{x}_i , the corresponding low-

dimensional representation \mathbf{z}_i consists of \mathbf{z}_i^1 and \mathbf{z}_i^2 . Note that \mathbf{z}_i^1 and \mathbf{z}_i^2 are further normalized to the unit sphere in \mathbb{R}^{d_p} , based on which we update class-specific prototypes $C \in \mathbb{R}^{K \times d_p}$ in a moving-average fashion. We call $f(\cdot)$ and $g(\cdot)$ as representor when the context is clear. Meanwhile, the classifier $h(\cdot)$ receives \mathbf{v}_i^1 as input and outputs $\mathbf{r}_i = h(\mathbf{v}_i^1)$. Besides the classification loss term implemented by a cross-entropy loss, we propose a prototypical alignment loss term which reduces the diversity between a similarity distribution over classes based on a softmax function over distances to class-specific prototypes and the probability disambiguated from the linear classifier prediction. Based on the feature encoder shared by linear classifier and projector, we build the connection between the effective representation space and the label disambiguation to facilitate the model training.

3.2. PLL with A Guided Prototypical Classifier

As we mentioned earlier, PaPi directly eliminates the contrastive learning module and adopts the opposite direction of disambiguation guidance compared to PiCO. Specifically, PaPi introduces a prototypical alignment loss term which guides the optimization of a prototypical classifier.

Prototypical alignment. To begin with, we describe prototypical alignment loss term. Given each sample (\mathbf{x}_i, Y_i) , we produce a prototypical similarity distribution \mathbf{s}_i^l over classes by a softmax function over distances to class-specific prototypes:

$$s_{ij}^l = \frac{\exp(\mathbf{z}_i^l \cdot \mathbf{c}_j / \tau)}{\sum_{k=1}^K \exp(\mathbf{z}_i^l \cdot \mathbf{c}_k / \tau)}, \quad (1)$$

where $l \in \{1, 2\}$ denotes different augmentation views, τ is the temperature and we use cosine similarity as distance measure. Meanwhile, we update the disambiguated probability based on the linear classifier prediction via putting more mass on more possible candidate labels:

$$u_{ij} = \begin{cases} \frac{r_{ij}}{\sum_{l \in Y_i} r_{il}} & \text{if } j \in Y_i, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

where j denotes the indices of candidate labels. Due to the low-quality prediction in initial training stage, we further adopt a soft and moving-average mechanism to update the probability, which also stabilizes the training process:

$$\mathbf{p}_i = \lambda \mathbf{p}_i + (1 - \lambda) \mathbf{u}_i, \quad (3)$$

where λ is a balancing factor. Intuitively speaking, initializing the disambiguated probability \mathbf{p}_i as a uniform distribution leads a good start as the representations are less distinguishable at the beginning. Therefore, we set $p_{ij} = 1/|Y_i|$

Algorithm 1 Pseudo-code of PaPi.

Input: Training dataset \mathcal{D} , encoder $f(\cdot)$, projector $g(\cdot)$, classifier $h(\cdot)$, mini-batch size B , epochs T_{max} , hyper-parameters $\lambda, \phi, \alpha, \gamma$.

- 1: **for** $t = 1, 2, \dots, T_{max}$ **do**
- 2: Shuffle \mathcal{D} into $\frac{|\mathcal{D}|}{B}$ mini-batches;
- 3: **for** $b = 1, 2, \dots, \frac{|\mathcal{D}|}{B}$ **do**
- 4: Obtain the disambiguation target by Eq.(3);
- 5: Create new training samples by Eq.(5);
- 6: Generate prototypical similarity distribution for these new samples by Eq.(1);
- 7: Update parameter of $f(\cdot)$, $g(\cdot)$ and $h(\cdot)$ by minimizing the empirical risk with loss in Eq.(9);
- 8: Update prototypes by Eq.(7);
- 9: **end for**
- 10: **end for**

Output: parameter of $f(\cdot)$ and $h(\cdot)$.

if $j \in Y_i$, otherwise $p_{ij} = 0$. Now we implement the per-sample alignment loss term by minimizing the Kullback-Leibler (KL) divergence of the prototypical similarity distribution \mathbf{s}_i^l and the disambiguated probability \mathbf{p}_i :

$$\mathcal{L}_{pa}^i(\mathbf{p}_i, \mathbf{s}_i^l) = \sum_{l \in \{1, 2\}} D_{KL}(\mathbf{p}_i || \mathbf{s}_i^l). \quad (4)$$

Considering that the label ambiguity generally exists in the training process, we further regularize the model from memorizing such ambiguous labels to learn more robust representations inspired by Mixup [39]. Specifically, we construct new training samples by linearly interpolating a sample (indexed by i) with another sample (indexed by $m(i)$) obtained from the same mini-batch randomly:

$$\hat{\mathbf{x}}_i = \phi \mathbf{x}_i + (1 - \phi) \mathbf{x}_{m(i)}, \quad (5)$$

where $\phi \sim Beta(\alpha, \alpha)$ and α is a hyperparameter. Denoting $\hat{\mathbf{s}}_i^l$ as the prototypical similarity distribution for $\hat{\mathbf{x}}_i$, we define the per-sample augmented alignment loss term as a weighted combination of two original loss term with respect to \mathbf{p}_i and $\mathbf{p}_{m(i)}$:

$$\mathcal{L}_{ali}^i = \phi \mathcal{L}_{pa}^i(\mathbf{p}_i, \hat{\mathbf{s}}_i^l) + (1 - \phi) \mathcal{L}_{pa}^i(\mathbf{p}_{m(i)}, \hat{\mathbf{s}}_i^l). \quad (6)$$

Prototype evolving. Instead of calculating the prototype representation as the normalized mean representation after every training iteration, we update the prototype \mathbf{c}_k of class k as a moving-average of the normalized representations for samples with pseudo-label k :

$$\mathbf{c}_k = \gamma \mathbf{c}_k + (1 - \gamma) \mathbf{z}_j, \quad \text{if } j \in I_k, \quad (7)$$

where I_k denotes the index set of samples pseudo-labeled with class k and γ is a balancing factor. Note that the prototypes are further normalized to the unit sphere.

Dataset	Method	$q = 0.1$	$q = 0.3$	$q = 0.5$	$q = 0.7$
Fashion-MNIST	Supervised	96.41 ± 0.02%			
	PaPi (ours)	95.79 ± 0.16%	95.63 ± 0.16%	95.60 ± 0.07%	95.16 ± 0.10%
	DPLL	95.76 ± 0.11%	95.61 ± 0.15%	95.52 ± 0.08%	94.95 ± 0.07%
	PiCO	94.89 ± 0.02%	94.35 ± 0.09%	93.87 ± 0.16%	90.05 ± 2.04%
	PRODEN	94.62 ± 0.25%	94.40 ± 0.08%	93.92 ± 0.27%	93.62 ± 0.28%
	LWS	94.73 ± 0.15%	94.50 ± 0.10%	93.16 ± 0.74%	91.08 ± 0.48%
	RC	95.27 ± 0.26%	95.02 ± 0.07%	94.73 ± 0.14%	93.92 ± 0.37%
	CC	95.30 ± 0.22%	94.95 ± 0.16%	94.64 ± 0.14%	93.42 ± 0.24%
SVHN	Supervised	97.64 ± 0.03%			
	PaPi (ours)	97.57 ± 0.04%	97.48 ± 0.07%	97.52 ± 0.04%	97.29 ± 0.07%
	DPLL	97.27 ± 0.06%	97.16 ± 0.07%	96.82 ± 0.08%	95.14 ± 0.10%
	PiCO	95.71 ± 0.12%	95.48 ± 0.16%	94.58 ± 0.33%	94.15 ± 0.12%
	PRODEN	96.82 ± 0.03%	96.61 ± 0.10%	96.60 ± 0.21%	95.49 ± 0.28%
	LWS	96.37 ± 0.01%	96.32 ± 0.43%	34.96 ± 0.15%	18.37 ± 1.12%
	RC	97.11 ± 0.10%	97.13 ± 0.04%	97.08 ± 0.06%	96.33 ± 0.05%
	CC	97.02 ± 0.03%	96.96 ± 0.05%	96.79 ± 0.10%	96.24 ± 0.12%
CIFAR-10	Supervised	97.36 ± 0.04%			
	PaPi (ours)	97.33 ± 0.06%	97.26 ± 0.08%	96.90 ± 0.09%	96.58 ± 0.07%
	DPLL	95.85 ± 0.22%	95.57 ± 0.15%	95.31 ± 0.11%	94.08 ± 0.16%
	PiCO	95.09 ± 0.31%	94.52 ± 0.39%	93.97 ± 0.55%	92.88 ± 0.57%
	PRODEN	93.85 ± 0.60%	93.36 ± 0.53%	92.95 ± 0.37%	90.98 ± 0.28%
	LWS	91.33 ± 0.08%	90.17 ± 0.24%	65.54 ± 1.64%	44.73 ± 1.43%
	RC	94.84 ± 0.38%	94.64 ± 0.07%	93.11 ± 0.03%	86.11 ± 2.85%
	CC	94.65 ± 0.27%	94.31 ± 0.20%	92.48 ± 0.33%	89.78 ± 0.35%
CIFAR-100-H	Supervised	82.31 ± 0.06%			
	PaPi (ours)	79.54 ± 0.13%	79.38 ± 0.19%	78.75 ± 0.17%	76.25 ± 0.24%
	DPLL	75.71 ± 0.12%	74.55 ± 0.09%	73.98 ± 0.12%	71.21 ± 0.18%
	PiCO	73.81 ± 0.13%	72.45 ± 0.44%	72.03 ± 0.19%	68.71 ± 1.20%
	PRODEN	73.24 ± 0.19%	72.21 ± 0.97%	70.11 ± 0.14%	67.25 ± 0.88%
	LWS	64.77 ± 1.45%	44.25 ± 0.68%	40.23 ± 1.27%	33.21 ± 1.22%
	RC	72.31 ± 0.16%	71.78 ± 0.12%	69.88 ± 0.24%	52.14 ± 1.22%
	CC	72.57 ± 0.88%	71.33 ± 0.90%	70.22 ± 0.87%	65.22 ± 0.89%

Table 1. Classification accuracy (mean ± std) on Fashion-MNIST, SVHN, CIFAR-10 and CIFAR-100-H with uniform partial labels under different ambiguity levels. The best accuracy is highlighted in bold.

Label disambiguation. As mentioned earlier, we build the connection between the highly effective representation space and the label disambiguation based on the weight-shared encoder. For multi-class classification, we adopt cross-entropy loss as the classification loss term with \mathbf{p}_i being the learning target. The per-sample classification loss term is defined as:

$$\mathcal{L}_{cla}^i = - \sum_{j=1}^K p_{ij} \cdot \log r_{ij}. \quad (8)$$

Model updating. Putting it all together, the overall training objective for each sample is to minimize a weighted sum

of all loss terms:

$$\mathcal{L}^i = \mathcal{L}_{cla}^i + \varphi(t) \cdot \mathcal{L}_{ali}^i, \quad (9)$$

where we employ a dynamic balancing function *w.r.t.* the epoch number t . We update the model by stochastic gradient descent without back-propagating towards the dashed arrow in Fig. 3. The pseudo-code is shown in Algorithm 1.

4. Experiments

4.1. Setup

Datasets. We adopt four widely used benchmark datasets including Fashion-MNIST [30], SVHN [17], CIFAR-10

Dataset	Method	$q = 0.01$	$q = 0.05$	$q = 0.1$	$q = 0.2$
CIFAR-100	Supervised	82.31 \pm 0.06%			
	PaPi (ours)	82.23 \pm 0.27%	81.60 \pm 0.18%	81.65 \pm 0.27%	79.49 \pm 0.22%
	DPLL	79.74 \pm 0.07%	78.97 \pm 0.13%	78.51 \pm 0.24%	75.77 \pm 0.06%
	PiCO	73.78 \pm 0.15%	72.78 \pm 0.38%	71.55 \pm 0.31%	48.76 \pm 1.19%
	PRODEN	73.20 \pm 0.18%	72.13 \pm 1.10%	71.88 \pm 0.32%	52.10 \pm 0.67%
	LWS	64.55 \pm 1.98%	50.19 \pm 0.34%	44.93 \pm 1.09%	37.97 \pm 1.42%
	RC	75.36 \pm 0.06%	74.44 \pm 0.31%	73.79 \pm 0.29%	64.74 \pm 0.82%
	CC	76.16 \pm 0.32%	75.04 \pm 0.10%	73.56 \pm 0.27%	71.43 \pm 0.32%
Mini-Imagenet	Supervised	74.45 \pm 0.18%			
	PaPi (ours)	72.43 \pm 0.02%	72.08 \pm 0.12%	71.54 \pm 0.34%	67.81 \pm 0.36%
	DPLL	72.20 \pm 0.01%	71.45 \pm 0.25%	70.89 \pm 0.28%	61.19 \pm 0.73%
	PiCO	67.24 \pm 1.73%	66.44 \pm 1.41%	64.00 \pm 1.62%	43.46 \pm 1.87%
	PRODEN	65.15 \pm 0.80%	61.70 \pm 0.68%	58.28 \pm 0.88%	37.04 \pm 0.72%
	LWS	65.85 \pm 0.51%	59.39 \pm 0.90%	54.12 \pm 0.25%	21.62 \pm 0.62%
	RC	64.60 \pm 0.55%	59.25 \pm 0.70%	51.11 \pm 0.11%	22.70 \pm 1.12%
	CC	64.74 \pm 0.44%	58.92 \pm 1.06%	49.23 \pm 0.90%	20.73 \pm 0.61%

Table 2. Classification accuracy (mean \pm std) on CIFAR-100 and Mini-Imagenet with uniform partial labels under different ambiguity levels. The best accuracy is highlighted in bold.

and CIFAR-100 [10]. We manually corrupt these datasets into partially labeled versions by uniform [24] and instance-dependent [33] generation process. We further investigate PaPi on fine-grained image classification dataset CIFAR-100-H [24] and few-shot image classification dataset Mini-Imagenet [21]. We set $q \in \{0.1, 0.3, 0.5, 0.7\}$ for Fashion-MNIST, SVHN, CIFAR-10 and CIFAR-100-H, set $q \in \{0.01, 0.05, 0.1, 0.2\}$ for CIFAR-100 and Mini-Imagenet.

Baselines. We compare PaPi against 7 state-of-the-art PLL methods: (1) DPLL [27] performs supervised learning on non-candidate labels and employs consistency regularization on candidate labels. (2) PiCO [24] identifies the true label by using contrastively learned representations. (3) VALEN [33] iteratively recovers the latent label distribution by adopting the variational inference. (4) PRODEN [15] accomplishes the update of model and the identification of true labels in a seamless manner. (5) LWS [26] considers the trade-off between losses on candidate labels and non-candidate labels. (6) RC [4] utilizes the importance reweighting strategy. (7) CC [4] employs a transition matrix to form an empirical risk estimator.

Implementation. For all datasets, we use an 18-layer ResNet following [7, 24] as the encoder $f(\cdot)$. We instantiate $g(\cdot)$ as a 2-layer MLP with one hidden layer (as well as ReLU activation). The classifier $h(\cdot)$ is a single linear layer, which is also used for inference. We present results based on three independent runs for all methods. Source code is available at <https://github.com/AlphaXia/PaPi>.

4.2. Main Results

PaPi achieves sota results. We present results on Fashion-MNIST, SVHN and CIFAR-10 in Table 1. Obviously, we can observe that PaPi significantly outperforms all baselines on all benchmark datasets, especially under *high ambiguity levels*. Specifically, PaPi outperforms current state-of-the-art method by **0.96%** and **2.50%** respectively on SVHN and CIFAR-10 when q is set to 0.7. More importantly, PaPi achieves comparable results against supervised learning model, showing that label disambiguation is satisfactorily completed in PaPi. For example, PaPi suffers tiny performance degradation (lower than 1%) compared with supervised learning model in many cases. The performance degradation of PaPi is lower than **3%** compared with supervised learning model *even* under $q = 0.7$, while most baselines demonstrate a significant performance drop. In Table 2, we also observe that PaPi consistently outperforms all baselines on CIFAR-100 and the superiority is more substantial than that on the dataset with a small label space. For example, PaPi outperforms the current state-of-the-art method by **3.14%** and **3.72%** respectively when q is set to 0.1 and 0.2.

In a more practical and challenge PLL setting where the partial labels are generated in an instance-dependent fashion, PaPi also outperforms the current state-of-the-art method by a large margin. As presented in Table 3, PaPi outperforms current state-of-the-art method by **4.40%** on CIFAR-10 and **4.57%** on CIFAR-100, which fully demonstrates the effectiveness of PaPi when facing more practical instance-dependent ambiguity.

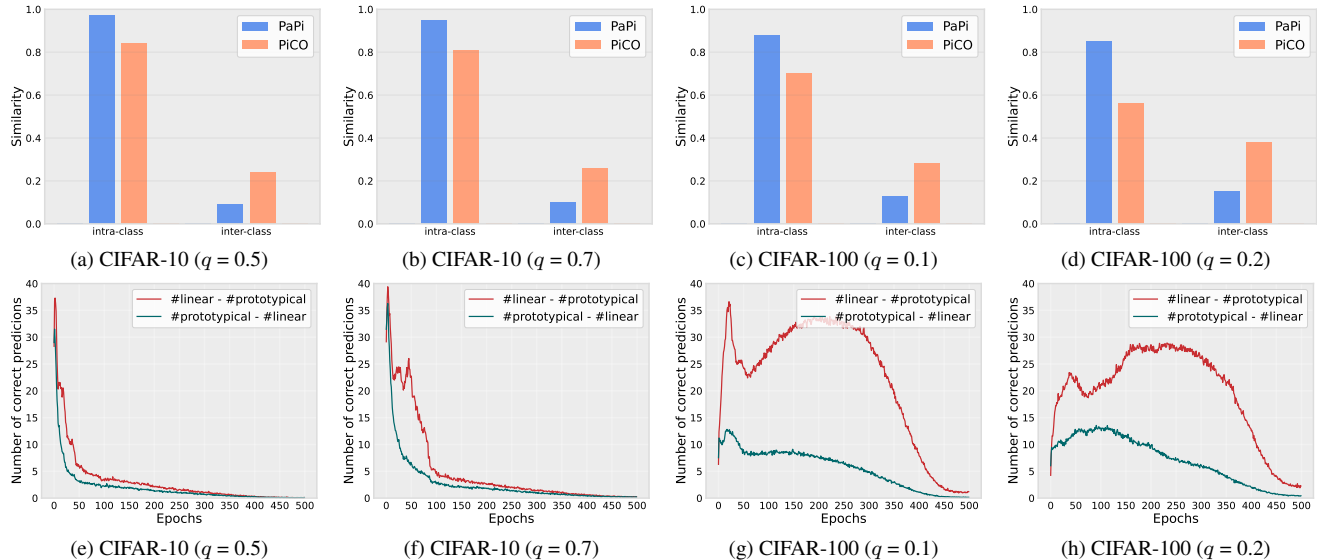


Figure 4. Visualizations of the effectiveness of learned representation. (a)-(d): Intra-class similarity and inter-class similarity on CIFAR-10 ($q = 0.5, 0.7$) and CIFAR-100 ($q = 0.1, 0.2$). (e)-(h): The red lines indicate the number of samples that were correctly classified by the linear classifier and incorrectly classified by the prototypical classifier per mini-batch, and the green lines are the opposite.

Method	Fashion-MNIST	SVHN	CIFAR-10	CIFAR-100
PaPi (ours)	93.84 ± 0.07%	97.74 ± 0.05%	95.13 ± 0.07%	63.70 ± 0.19%
DPLL	92.72 ± 0.14%	92.11 ± 0.88%	90.73 ± 0.11%	33.14 ± 1.74%
PiCO	87.86 ± 0.60%	91.18 ± 1.10%	90.50 ± 0.17%	59.13 ± 0.31%
VALEN	88.89 ± 0.13%	94.66 ± 0.17%	90.16 ± 0.52%	29.25 ± 0.14%
PRODEN	91.52 ± 0.24%	95.78 ± 0.46%	89.10 ± 0.44%	31.25 ± 0.20%
LWS	71.68 ± 3.89%	54.49 ± 7.21%	48.67 ± 7.68%	25.14 ± 2.18%
RC	91.42 ± 0.30%	94.85 ± 0.26%	88.66 ± 0.37%	38.25 ± 0.79%
CC	91.43 ± 0.76%	95.08 ± 0.24%	88.33 ± 0.75%	43.30 ± 1.90%

Table 3. Classification accuracy (mean ± std) on Fashion-MNIST, SVHN, CIFAR-10 and CIFAR-100 with instance-dependent partial labels. The best accuracy is highlighted in bold.

From the above experimental results, we show that PaPi can effectively handle both uniform and instance-dependent PLL setting. Besides, we investigate the performance of PaPi and all baselines on CIFAR-100-H and Mini-Imagenet, which are less explored in previous literature. From Table 1 and Table 2, we observe that PaPi still outperforms all baselines significantly. The whole comparison results emphasize the superiority of our framework.

PaPi learns effective representations. We validate the effectiveness of learned representation by calculating the intra-class and inter-class similarity. From Fig. 4a to Fig. 4d, we can observe that PaPi produces well-separated clusters and more distinguishable representations. Moreover, we justify our correct disambiguation guidance direction. From Fig. 4e to Fig. 4h, we find that the linear classi-

fier always shows higher accuracy than the prototypical one and the linear classifier always has something new to teach the prototypical one until convergence.

4.3. Ablation and Analysis

The effectiveness of prototypical alignment term. We explore the effectiveness of our proposed prototypical alignment. Specifically, we compare PaPi with several variants: (1) *Variant 1* which removes the mixup augmentation compared with PaPi; (2) *Variant 2* which removes the prototypical alignment loss compared with PaPi. From Table 4, we observe that PaPi outperforms *Variant 1* (e.g., +1.47% on CIFAR-100 ($q = 0.2$)), which verifies the effectiveness of mixup augmentation. We also observe that PaPi outperforms *Variant 2* (e.g., +7.19% on CIFAR-100 ($q = 0.2$)), which confirms the importance of prototypical alignment

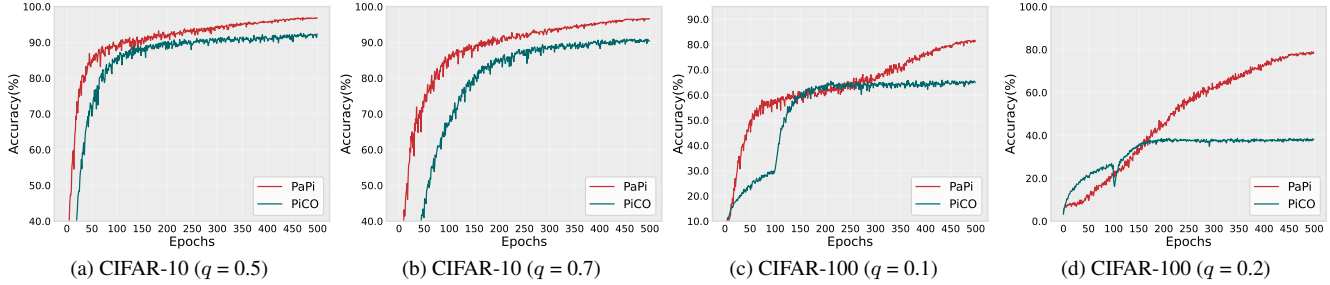


Figure 5. Classification accuracy calculated with the learned prototypes on CIFAR-10 ($q = 0.5, 0.7$) and CIFAR-100 ($q = 0.1, 0.2$).

Method	CIFAR-10		CIFAR-100	
	$q = 0.5$	$q = 0.7$	$q = 0.1$	$q = 0.2$
PaPi (ours)	96.90 ± 0.09%	96.58 ± 0.07%	81.65 ± 0.27%	79.49 ± 0.22%
Variant 1	96.42 ± 0.11%	95.82 ± 0.12%	79.82 ± 0.09%	78.02 ± 0.05%
Variant 2	94.39 ± 0.10%	92.26 ± 0.12%	77.48 ± 0.07%	72.30 ± 0.26%
PiCO	93.97 ± 0.55%	92.88 ± 0.57%	71.55 ± 0.31%	48.76 ± 1.19%

Table 4. Classification accuracy (mean ± std) of different degenerated methods on CIFAR-10 ($q = 0.5, 0.7$) and CIFAR-100 ($q = 0.1, 0.2$).

Data augmentation compositions	CIFAR-10 $q = 0.5$	CIFAR-100 $q = 0.1$
PaPi (S+W)	96.90 ± 0.09%	81.65 ± 0.27%
PaPi (2×S)	96.91 ± 0.05%	80.96 ± 0.08%
PaPi (2×W)	96.60 ± 0.03%	79.35 ± 0.14%
DPLL (2S+W)	95.31 ± 0.11%	78.51 ± 0.24%
DPLL (3×S)	95.67 ± 0.06%	78.41 ± 0.50%
DPLL (3×W)	93.76 ± 0.21%	75.23 ± 0.24%

Table 5. Accuracy (mean ± std) under different data augmentation compositions on CIFAR-10 ($q = 0.5$) and CIFAR-100 ($q = 0.1$).

for identifying the true label. Besides, we find that even *Variant 2* achieves comparable results against the current state-of-the-art method, which validates the importance of self-teaching fashion.

The performance of learned prototypical classifier. In Fig. 5, we report the classification accuracy calculated with the learned prototypes. We can observe that PaPi achieves better performance especially facing *high ambiguity levels*, which demonstrates the effectiveness of our learned prototypical classifier. For example, PaPi outperforms PiCO by **6.24%** and **40.29%** respectively on CIFAR-10 ($q = 0.7$) and CIFAR-100 ($q = 0.2$).

The impact of data augmentation compositions. In our main experiments, we adopt one weak and one strong augmentation. As shown in Table 5, we compare the accuracy

when equipped with different data augmentation compositions, where S means strong augmentation and W means weak augmentation. We observe that the combination of strong and weak augmentation achieves the best performance. Similarly, we evaluate DPLL [27] under different compositions. When equipped with two weak augmentations, PaPi suffers less performance drop compared with DPLL, which demonstrates the competitiveness of PaPi.

5. Conclusion

In this paper, we proposed a simple PLL framework termed PaPi which explicitly contrasts the prototypical similarity with the visual similarity between categories, such that PaPi is remarkable for improving the class-level discrimination of learned representation. Extensive experimental results under multiple PLL settings demonstrated PaPi established new state-of-the-art performance especially when facing high ambiguity levels and instance-dependent ambiguity.

Acknowledgement. This research was supported by the National Key Research & Development Plan of China (No. 2018AAA0100104), the National Science Foundation of China (62125602, 62076063, 62206050), China Post-doctoral Science Foundation under Grant 2021M700023, Jiangsu Province Science Foundation for Youths under Grant BK20210220, the Young Elite Scientists Sponsorship Program of Jiangsu Association for Science and Technology TJ-2022-078, the Big Data Computing Center of Southeast University.

References

- [1] C. Chen, V. M. Patel, and R. Chellappa. Learning from ambiguously labeled face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(7):1653–1667, 2017. [1](#)
- [2] T. Chen, S. Kornblith, M. Norouzi, and G. E. Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of 37th International Conference on Machine Learning (ICML'20)*, pages 1597–1607, 2020. [3](#)
- [3] T. Cour, B. Sapp, and B. Taskar. Learning from partial labels. *Journal of Machine Learning Research*, 12(5):1501–1536, 2011. [1, 2](#)
- [4] L. Feng, J. Lv, B. Han, M. Xu, G. Niu, X. Geng, B. An, and M. Sugiyama. Provably consistent partial-label learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS'20)*, pages 10948–10960, 2020. [1, 3, 6](#)
- [5] J. Grill, F. Strub, F. Althé, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. Ávila Pires, Z. Guo, M. Gheshlaghi Azar, B. Piot, K. Kavukcuoglu, R. Munos, and M. Valko. Bootstrap your own latent - a new approach to self-supervised learning. *Advances in Neural Information Processing Systems 33 (NeurIPS'20)*, pages 21271–21284, 2020. [3](#)
- [6] T. Han, J. Gao, Y. Yuan, and Q. Wang. Unsupervised semantic aggregation and deformable template matching for semi-supervised learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS'20)*, 2020. [3](#)
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of 29th IEEE conference on Computer Vision and Pattern Recognition (CVPR'16)*, pages 770–778, 2016. [6](#)
- [8] E. Hüllermeier and J. Beringer. Learning from ambiguously labeled examples. *Intelligent Data Analysis*, 10(5):419–439, 2006. [1, 2](#)
- [9] R. Jin and Z. Ghahramani. Learning with multiple labels. In *Advances in Neural Information Processing Systems 16 (NeurIPS'03)*, pages 921–928, 2003. [1, 2](#)
- [10] A. Krizhevsky and G. Hinton. *Learning multiple layers of features from tiny images*. Citeseer, 2009. [6](#)
- [11] J. Li, C. Xiong, and S. C. H. Hoi. Mopro: Webly supervised learning with momentum prototypes. In *Proceedings of 9th International Conference on Learning Representations (ICLR'21)*, 2021. [3](#)
- [12] J. Li, P. Zhou, C. Xiong, and S. C. H. Hoi. Prototypical contrastive learning of unsupervised representations. In *Proceedings of 9th International Conference on Learning Representations (ICLR'21)*, 2021. [3](#)
- [13] J. Luo and F. Orabona. Learning from candidate labeling sets. In *Advances in Neural Information Processing Systems 23 (NeurIPS'10)*, pages 1504–1512, 2010. [1](#)
- [14] J. Lv, B. Liu, L. Feng, N. Xu, M. Xu, B. An, G. Niu, X. Geng, and M. Sugiyama. On the robustness of average losses for partial-label learning. *arXiv preprint arXiv:2106.06152*, 2021. [3](#)
- [15] J. Lv, M. Xu, L. Feng, G. Niu, X. Geng, and M. Sugiyama. Progressive identification of true labels for partial-label learning. In *Proceedings of 37th International Conference on Machine Learning (ICML'20)*, pages 6500–6510, 2020. [1, 3, 6](#)
- [16] G. Lyu, S. Feng, T. Wang, C. Lang, and Y. Li. Gm-pll: Graph matching based partial label learning. *IEEE Transactions on Knowledge and Data Engineering*, 33(2):521–535, 2021. [2](#)
- [17] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng. Reading digits in natural images with unsupervised feature learning. *Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. [5](#)
- [18] F. Pahde, M. Puszcz, T. Klein, and M. Nabi. Multimodal prototypical networks for few-shot learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2644–2653, 2021. [3](#)
- [19] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems 30 (NeurIPS'17)*, pages 4077–4087, 2017. [1, 3](#)
- [20] C. Tang and M. Zhang. Confidence-rated discriminative partial label learning. In *Proceedings of 31st AAAI Conference on Artificial Intelligence (AAAI'17)*, pages 2611–2617, 2017. [1, 2](#)
- [21] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. In *Advances in neural information processing systems 29 (NeurIPS'16)*, volume 29, 2016. [6](#)
- [22] H. Wang, W. Liu, Y. Zhao, T. Hu, K. Chen, and G. Chen. Learning from multi-dimensional partial labels. In *Proceedings of 29th International Joint Conference on Artificial Intelligence (IJCAI'20)*, pages 2943–2949, 2020. [1](#)
- [23] H. Wang, Y. Qiang, C. Chen, W. Liu, T. Hu, Z. Li, and G. Chen. Online partial label learning. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, pages 455–470, 2020. [1](#)
- [24] H. Wang, R. Xiao, Y. Li, L. Feng, G. Niu, G. Chen, and J. Zhao. Pico: Contrastive label disambiguation for partial label learning. In *Proceedings of 10th International Conference on Learning Representations (ICLR'22)*, 2022. [1, 3, 6](#)
- [25] T. Wei, J. Shi, Y. Li, and M. Zhang. Prototypical classifier for robust class-imbalanced learning. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 44–57, 2022. [3](#)
- [26] H. Wen, J. Cui, H. Hang, J. Liu, Y. Wang, and Z. Lin. Leveraged weighted loss for partial label learning. In *Proceedings of 36th International Conference on Machine Learning (ICML'21)*, pages 11091–11100, 2021. [1, 6](#)
- [27] D. Wu, D. Wang, and M. Zhang. Revisiting consistency regularization for deep partial label learning. In *Proceedings of 37th International Conference on Machine Learning (ICML'22)*, volume 162, pages 24212–24225, 2022. [1, 3, 6, 8](#)
- [28] Z. Wu, J. Lv, and M. Sugiyama. Learning with proper partial labels. *Neural Computation*, 35(1):58–81, 2023. [1, 3](#)
- [29] S. Xia, J. Lv, N. Xu, and X. Geng. Ambiguity-induced contrastive learning for instance-dependent partial label learning.

- ing. In *Proceedings of 31st International Joint Conference on Artificial Intelligence (IJCAI'22)*, pages 3615–3621, 2022. 3
- [30] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. 5
- [31] N. Xu, Y. Liu, and X. Geng. Label enhancement for label distribution learning. *IEEE Transactions on Knowledge and Data Engineering*, 33(4):1632–1643, 2019. 3
- [32] N. Xu, J. Lv, and X. Geng. Partial label learning via label enhancement. In *Proceedings of 33rd AAAI Conference on Artificial Intelligence (AAAI'19)*, pages 5557–5564, 2019. 1, 2
- [33] N. Xu, C. Qiao, X. Geng, and M. Zhang. Instance-dependent partial label learning. In *Advances in Neural Information Processing Systems 34 (NeurIPS'21)*, 2021. 3, 6
- [34] N. Xu, J. Shu, Y. Liu, and X. Geng. Variational label enhancement. In *Proceedings of 37th International Conference on Machine Learning (ICML'20)*, pages 10597–10606, 2020. 3
- [35] W. Xu, Y. Xian, J. Wang, B. Schiele, and Z. Akata. Attribute prototype network for zero-shot learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS'20)*, 2020. 3
- [36] Y. Yao, C. Gong, J. Deng, X. Chen, J. Wu, and J. Yang. Deep discriminative cnn with temporal ensembling for ambiguously-labeled image classification. In *Proceedings of 34th AAAI Conference on Artificial Intelligence (AAAI'20)*, pages 12669–12676, 2020. 1, 3
- [37] Y. Yao, C. Gong, J. Deng, and J. Yang. Network cooperation with progressive disambiguation for partial label learning. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, pages 471–488, 2020. 1, 3
- [38] F. Yu and M. Zhang. Maximum margin partial label learning. *Machine Learning*, 106(4):573–593, 2017. 1, 2
- [39] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. mixup: Beyond empirical risk minimization. In *Proceedings of 6th International Conference on Learning Representations (ICLR'18)*, 2018. 4
- [40] M. Zhang and F. Yu. Solving the partial label learning problem: An instance-based approach. In *Proceedings of 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, pages 4048–4054, 2015. 1, 2
- [41] M. Zhang, B. Zhou, and X. Liu. Partial label learning via feature-aware disambiguation. In *Proceedings of 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'16)*, pages 1335–1344, 2016. 1, 2