

# Learning with Partial Labels from Semi-supervised Perspective

Ximing Li<sup>1,2</sup>, Yuanzhi Jiang<sup>1,2</sup>, Changchun Li<sup>1,2,\*</sup>, Yiyuan Wang<sup>3,4</sup>, Jihong Ouyang<sup>1,2</sup>

<sup>1</sup>College of Computer Science and Technology, Jilin University, China

<sup>2</sup>Key Laboratory of Symbolic Computation and Knowledge Engineering of MOE, Jilin University, China

<sup>3</sup>College of Information Science and Technology, Northeast Normal University, China

<sup>4</sup>Key Laboratory of Applied Statistics of MOE, Northeast Normal University, China

liximing86@gmail.com, yzjiang20@mails.jlu.edu.cn, changchunli93@gmail.com, wangyy912@nenu.edu.cn, ouyj@jlu.edu.cn


## Abstract

Partial Label (PL) learning refers to the task of learning from the partially labeled data, where each training instance is ambiguously equipped with a set of candidate labels but only one is valid. Advances in the recent deep PL learning literature have shown that the deep learning paradigms, *e.g.*, self-training, contrastive learning, or class activate values, can achieve promising performance. Inspired by the impressive success of deep Semi-Supervised (SS) learning, we transform the PL learning problem into the SS learning problem, and propose a novel PL learning method, namely Partial Label learning with Semi-supervised Perspective (PLSP). Specifically, we first form the pseudo-labeled dataset by selecting a small number of reliable pseudo-labeled instances with high-confidence prediction scores and treating the remaining instances as pseudo-unlabeled ones. Then we design a SS learning objective, consisting of a supervised loss for pseudo-labeled instances and a semantic consistency regularization for pseudo-unlabeled instances. We further introduce a complementary regularization for those non-candidate labels to constrain the model predictions on them to be as small as possible. Empirical results demonstrate that PLSP significantly outperforms the existing PL baseline methods, especially on high ambiguity levels. Code available: <https://github.com/changchunli/PLSP>.

## Introduction

During the past decades, modern deep neural networks have gained great success in various domains such as computer vision and natural language processing. Commonly, they are built on the paradigm of supervised learning, which often requires massive training instances with precise labels. However, in many real-world scenarios, the high-quality training instances are intractable to collect, because instance annotation by human-beings is costly and even subject to label ambiguity and noise, potentially resulting in many training data with various noisy supervision (Li, Socher, and Hoi 2020; Li et al. 2022). Among them, one prevalent noisy challenge is from the partially labeled data, where each training instance is equipped with a set of candidate labels but only one is valid (Cour, Sapp, and Taskar 2011). As illustrated in Fig.1, for a human annotator it could be difficult to correctly distinguish *Alaskan Malamute* and *Huskie*, so she/he

\*Corresponding author.



Annotator 1 : Huskie

Annotator 2 : Alaskan Malamute

Figure 1: An example of PL instances. An *Alaskan Malamute* is in this image, but annotators may also tag it with *Huskie*.

may tend to retain both of them as candidate labels. Due to the popularity of such noisy data in applications, *e.g.*, web mining (Luo and Orabona 2010), multimedia context analysis (Zeng et al. 2013), and image classification (Chen, Patel, and Chellappa 2018), the paradigm of learning from partial labels, formally dubbed as **Partial Label (PL)** learning, has recently attracted more attention from the machine learning community (Feng and An 2019b; Feng et al. 2020; Lv et al. 2020; Li, Li, and Ouyang 2020; Li et al. 2021; Wang et al. 2022a; Wu, Wang, and Zhang 2022).

Naturally, the main challenge of PL learning lies in the ambiguity of partial labels, because the ground-truth label is unknown and can not be directly accessible to the learning method. Accordingly, the mainstream of PL learning methods concentrates on recovering precise supervised signals from the ambiguous candidate labels. Some two-stage methods refine the candidate labels by label propagation among instance nearest neighbors (Zhang and Yu 2015; Zhang, Zhou, and Liu 2016; Xu, Lv, and Geng 2019); and most PL learning methods jointly train the classifier with the refined labels and refine the candidate labels with the classifier predictions (Wu and Zhang 2018; Zhang, Zhou, and Liu 2019; Feng and An 2019a; Li, Li, and Ouyang 2020; Ni et al. 2021). Besides them, some deep PL learning methods employ discriminators to recover precise supervision from the candidate labels under the frameworks of GAN (Zhang et al. 2020) and Triple-

GAN (Li et al. 2021). Despite promising performance, the refined labels of these PL learning methods can be still ambiguous and inaccurate for most training instances, resulting in potential performance degradation.

In parallel with PL learning, **Semi-Supervised (SS)** learning has recently achieved great progress with strong deep neural backbones (Berthelot et al. 2019, 2020; Xie et al. 2020; Sohn et al. 2020; Zhang et al. 2021; Li, Li, and Ouyang 2021). The recent deep SS learning methods are mainly based on the consistency regularization with the assumption that the classifier tends to give consistent predictions on augmented variants of the same instance. Inspired by them, we revisit the problem that the refined labels of PL learning methods may still be ambiguous for most training instances, and throw the following question: **“Whether we can efficiently select a small number of reliable instances with high-confidence pseudo-labels, and then resolve the PL learning task with the strong SS learning techniques?”**

Motivated by this question, we develop a novel PL learning method, namely **Partial Label learning with Semi-supervised Perspective (PLSP)**. Our PLSP consists of two stages. In the first stage, we efficiently per-train the classifier by treating all candidate labels equally important, and then select a small number of “reliable” pseudo-labeled instances with high-confidence predictions and treat the remaining instances as pseudo-unlabeled instances. In the second stage, we formulate a SS learning objective to induce the classifier over the pseudo-training dataset. To be specific, we incorporate a consistency regularization with respect to the weakly- and strongly-augmented instances, and draw semantic-transformed features for them to further achieve consistency at the semantic level. To efficiently optimize the objective with semantic-transformed features, we derive an approximation of its expectation form. We conduct extensive experiments to evaluate the effectiveness of PLSP on benchmark datasets. Empirical results demonstrate that PLSP is superior to the existing PL learning baseline methods, especially on high ambiguity levels.

## Related Work

### Partial Label Learning

There are many PL learning studies based on the shallow frameworks. The early disambiguation-free methods efficiently induce the classifiers by treating all candidate labels equally important (Cour et al. 2009; Cour, Sapp, and Taskar 2011; Zhang, Yu, and Tang 2017), while in PLSP we have also employed this spirit to per-train the classifier to initialize the pseudo-training dataset. Beyond them, the disambiguation methods aim to induce stronger classifiers by refining precise supervision from candidate labels (Wu and Zhang 2018; Feng and An 2018). Some two-stage methods first refine the candidate labels by label propagation among instance nearest neighbors (Zhang and Yu 2015; Zhang, Zhou, and Liu 2016; Xu, Lv, and Geng 2019). But most disambiguation methods jointly train the classifier with the refined labels and refine the candidate labels with the classifier predictions (Feng and An 2019a; Li, Li, and Ouyang 2020). However, the refined labels may be also noisy for most training instances.

Inspired by the effectiveness of deep learning and efficiency of stochastic optimization, a number of deep PL learning methods have been recently developed (Zhang et al. 2020; Lv et al. 2020; Feng et al. 2020; Wen et al. 2021; Yan and Guo 2020; Li et al. 2021; Xu et al. 2021; Zhang et al. 2022; Wang et al. 2022a). Some deep PL learning methods refine the candidate labels with adversarial training, such as the attempts based on GAN (Zhang et al. 2020) and Triple-GAN (Li et al. 2021). Most other methods design proper objectives for PL learning. For example, PRODEN (Lv et al. 2020) optimizes a classifier-consistent objective derived by the assumption that the ground-truth label would achieve the minimal loss among candidate labels; Feng et al. (2020) propose risk-consistent and classifier-consistent methods with the assumption that the candidate labels are drawn from a uniform distribution; and Wen et al. (2021) propose a risk-consistent leveraged weighted loss with label-specific candidate label sampling. However, those methods highly rely on their prior assumptions. Besides, the recent deep PL learning method PiCO (Wang et al. 2022a) borrows the idea of contrastive learning to keep the consistence between the augmented versions of each instance. In contrast to PiCO, we also employ a consistency regularization with respect to augmented instances but we further draw semantic-transformed features for them to achieve a new semantic consistency regularization.

### Semi-Supervised Learning

The recent deep SS learning methods are mainly based on the consistency regularization (Laine and Aila 2017; Tarvainen and Valpola 2017; Miyato et al. 2019; Berthelot et al. 2019, 2020; Xie et al. 2020; Sohn et al. 2020; Zhang et al. 2021). It is built on a simple concept that the classifier should give consistent predictions on an unlabeled instance and its perturbed version. To conduct this idea, many perturbation methods are adopted, such as the virtual adversarial training used in VAT (Miyato et al. 2019) and the mixup technique adopted by MixMatch (Berthelot et al. 2019). With the data augmentation technique popular, the previous arts (Berthelot et al. 2020; Xie et al. 2020; Sohn et al. 2020; Zhang et al. 2021) keep the classifier predictions on the weakly- and strongly-augmented variants of an unlabeled instance to be consistent, and empirically achieve impressive performance. The main data augmentation techniques used in these methods are at the pixel-level, such as flip-and-shift, Cutout (Devries and Taylor 2017), AutoAugment (Cubuk et al. 2019), CTAugment (Berthelot et al. 2020), and RandAugment (Cubuk et al. 2020) *etc.* Complementary to these pixel-level augmentations, Wang et al. (2022b) design a semantic-level data augmentation method motivated by the linear characteristic of deep features. In PLSP, we employ both pixel-level and semantic-level data augmentations, and design a semantic consistency regularization for PL learning by performing semantic-level transformations on both weakly- and strongly-augmented variants of an instance.

### The Proposed PLSP Approach

In this section, we introduce the proposed PL learning method, namely **Partial Label learning with Semi-supervised Perspective (PLSP)**.

**Problem formulation of PL learning** We now formulate the problem of PL learning. Let  $\mathcal{X} \subseteq \mathbb{R}^d$  be a  $d$ -dimensional feature space and  $\mathcal{Y} = \{1, \dots, l\}$ ,  $l \geq 2$  be the label space. In the context of PL learning, we are given by a training dataset consisting of  $n$  instances, denoted by  $\Omega = \{(\mathbf{x}_i, C_i)\}_{i=1}^n$ . For each instance,  $\mathbf{x}_i \in \mathcal{X}$  and  $C_i \subseteq \mathcal{C}$  are its feature vector and corresponding candidate label set, respectively, where  $\mathcal{C} = \{\mathcal{P}(\mathcal{Y}) \setminus \emptyset \setminus \mathcal{Y}\}$  is the power set of  $\mathcal{Y}$  except for the empty set and the whole label set. Specially, the single ground-truth label of each instance is unknown and must be concealed in its corresponding candidate label set. The goal of PL learning is to train a classifier  $f(\cdot; \Theta)$ , parameterized by  $\Theta$ , from such noisy training dataset  $\Omega$ .

### Overview of PLSP

The main idea of PLSP is to transform the PL learning problem into the SS learning problem, and then induce the classifier by leveraging the well-established SS learning paradigms. Specifically, we first select a small number of “reliable” partial-labeled instances (*i.e.*,  $m \ll n$ ) from  $\Omega$  according to their predicted scores, *e.g.*, class activation values (Zhang et al. 2022), and form a pseudo-labeled instance set  $\Omega_l = \{(\mathbf{x}_{p(i)}, y_{p(i)}) \in \mathcal{X} \times \mathcal{C}_{p(i)}\}_{i=1}^m$ , where the subscript  $p(i)$  denotes the mapping function of instance index and  $y_{p(i)}$  is the corresponding high-confidence pseudo label. We treat the remaining instances as pseudo-unlabeled instances, denoted by  $\Omega_u$ . Accordingly, we can further treat the pseudo-training dataset  $\{\Omega_l, \Omega_u\}$  as a training dataset of SS learning, so as to train a classifier from it by leveraging the following well-established objective of SS learning:

$$\mathcal{L}(\{\Omega_l, \Omega_u\}; \Theta) = \mathcal{L}_l(\Omega_l; \Theta) + \mathcal{R}_u(\Omega_u; \Theta), \quad (1)$$

where  $\mathcal{L}_l$  is the pseudo-supervised loss with respect to  $\Omega_l$ ; and  $\mathcal{R}_u$  the regularization with respect to  $\Omega_u$  (*e.g.*, consistency regularization). In the following, we will introduce the details of forming the pseudo-training dataset  $\{\Omega_l, \Omega_u\}$  and constructing SS learning loss over  $\{\Omega_l, \Omega_u\}$  for PL learning, and then show the specific objective of PLSP as well as the full training process.

### Forming the Pseudo-Training Dataset $\{\Omega_l, \Omega_u\}$

To form the pseudo-labeled instance set  $\Omega_l$ , we pre-train the classifier  $f(\cdot; \Theta)$  by using a simple disambiguation-free PL learning loss, where all candidate labels are treated equally:

$$\mathcal{L}_{df}(\Omega; \Theta) = \frac{1}{|\Omega|} \sum_{(\mathbf{x}_i, C_i) \in \Omega} \frac{1}{|C_i|} \sum_{j \in C_i} -\log p_{ij} \quad (2)$$

where  $\mathbf{p}_i = [p_{ij}]_{j \in \mathcal{Y}}^\top$  is the classifier prediction of instance  $\mathbf{x}_i$ , and  $p_{ij} = e^{z_{ij}} / \sum_{j' \in \mathcal{Y}} e^{z_{ij'}}$ ,  $\mathbf{z}_i = f(\mathbf{x}_i; \Theta)$ . With the per-trained classifier  $f(\cdot; \Theta)$ ,<sup>1</sup> we select a small number of “reliable” pseudo-labeled instances to form  $\Omega_l$ . Specifically, for each instance  $(\mathbf{x}_i, C_i) \in \Omega$ , we assign the candidate label with the highest class activation value (CAV) (Zhang et al. 2022) as its pseudo label:

$$y_i = \arg \max_{j \in C_i} v_{ij}, \quad v_{ij} = \tilde{z}_{ij} |\tilde{z}_{ij} - 1|, \quad \tilde{\mathbf{z}}_i = f(\mathbf{x}_i; \tilde{\Theta}).$$

<sup>1</sup>This pre-training process can be very efficient and converge within a few epochs.

For each class  $j \in \mathcal{Y}$ , we construct its pseudo-labeled instance set  $\Omega_l^j$  by choosing instances with the top- $k$  CAVs of class  $j$ :

$$\Omega_l^j = \{(\mathbf{x}_i, y_i) | i \in \text{TopK}(\{v_{ij} | (\mathbf{x}_i, C_i) \in \Omega, y_i = j\})\},$$

where, as its name suggests,  $\text{TopK}(\cdot)$  outputs the index set of instances with the top- $k$  CAVs. Accordingly, the pseudo-labeled set  $\Omega_l$  can be formed as follows:<sup>2</sup>

$$\Omega_l = \bigcup_{j \in \mathcal{Y}} \Omega_l^j, \quad (3)$$

and the remaining instances can constitute the pseudo-unlabeled instance set  $\Omega_u$  as follows:

$$\Omega_u = \{(\mathbf{x}_i, C_i) | (\mathbf{x}_i, C_i) \in \Omega, (\mathbf{x}_i, y_i) \notin \Omega_l\}. \quad (4)$$

### Forming the SS Learning Loss over $\{\Omega_l, \Omega_u\}$

Given  $\{\Omega_l, \Omega_u\}$ , we continue to optimize the per-trained classifier  $f(\cdot; \Theta)$  by using the SS learning loss of PLSP, including the pseudo-supervised loss  $\mathcal{L}_l$  for  $\Omega_l$  and the regularization term  $\mathcal{R}_u$  for  $\Omega_u$ .

**Pseudo-supervised loss.** We can treat  $\Omega_l$  as a labeled dataset, and directly formulate the specific pseudo-supervised loss as follows:

$$\mathcal{L}_l(\Omega_l; \Theta) = \frac{1}{|\Omega_l|} \sum_{(\mathbf{x}_i, y_i) \in \Omega_l} -\log p_{iy_i}. \quad (5)$$

**Regularizing the pseudo-unlabeled instances** Inspired by the impressive success of the consistency regularization in SS learning (Xie et al. 2020; Sohn et al. 2020; Zhang et al. 2021), we employ it to regularize the pseudo-unlabeled instances. Specifically, for each instance within  $\Omega_u$ , we first generate its weakly- and strongly-augmented variants with the wide-used pixel-level data augmentation tricks,<sup>3</sup> and then constrain their corresponding prediction scores to be consistent. Formally, for each pseudo-unlabeled instance  $(\mathbf{x}_i, C_i) \in \Omega_u$ , let its weakly- and strongly-augmented variants denote by  $\mathbf{x}_i^w = \alpha(\mathbf{x}_i)$  and  $\mathbf{x}_i^s = \mathcal{A}(\mathbf{x}_i)$ , respectively. Its corresponding consistency regularization term can be written as follows:

$$\mathfrak{R}_u((\mathbf{x}_i, C_i); \Theta) = h(\hat{\mathbf{p}}_i^w) \text{KL}(\hat{\mathbf{p}}_i^w || \mathbf{p}_i^s), \quad (6)$$

where  $\text{KL}(\cdot || \cdot)$  denotes the KL-divergence. More specially,  $\hat{\mathbf{p}}_i = [\hat{p}_{ij}]_{j \in \mathcal{Y}}^\top$  is the pseudo-target approximated on the weakly-augmented variant  $\mathbf{x}_i^w$ :

$$\hat{p}_{ij} = \frac{\mathbf{1}(j \in C_i) \hat{p}_{ij}^w}{\sum_{j' \in \mathcal{Y}} \mathbf{1}(j' \in C_i) \hat{p}_{ij'}^w},$$

$$\hat{p}_{ij}^w = \frac{e^{\tilde{z}_{ij}^w}}{\sum_{j' \in \mathcal{Y}} e^{\tilde{z}_{ij'}^w}}, \quad \tilde{\mathbf{z}}_i^w = f(\mathbf{x}_i^w; \hat{\Theta}),$$

<sup>2</sup>The total number of pseudo-labeled instances  $m = k \times l$ .

<sup>3</sup>These data augmentations include flip-and-shift, Cutout (Devries and Taylor 2017), AutoAugment (Cubuk et al. 2019), CTAugment (Berthelot et al. 2020), and RandAugment (Cubuk et al. 2020) *etc.* We will introduce their implementation details in the experiment part.

where  $\hat{\Theta}$  is the fixed copy of the current parameters  $\Theta$ ;  $\mathbf{p}_i^s = [p_{ij}^s]_{j \in \mathcal{Y}}^\top$  is the classifier prediction on the strongly-augmented variant  $\mathbf{x}_i^s$ , and  $p_{ij}^s = e^{z_{ij}^s} / \sum_{j' \in \mathcal{Y}} e^{z_{ij'}^s}$ ,  $\mathbf{z}_i^s = f(\mathbf{x}_i^s; \Theta)$ ;  $h(\hat{\mathbf{p}}_i^w)$  is an indicator function used to retain high-confidence pseudo-unlabeled instances in this consistency regularization term, specifically defined as follows:

$$h(\hat{\mathbf{p}}_i^w) = \mathbb{1} \left( \left( \max_{j \in \mathcal{Y}} \hat{p}_{ij}^w \geq \tau \right) \wedge \left( \arg \max_{j \in \mathcal{Y}} \hat{p}_{ij}^w \in C_i \right) \right),$$

where  $\tau \in (0.5, 1.0]$  is the confidence threshold. Accordingly, the overall consistency regularization over  $\Omega_u$  is stated as:

$$\mathcal{R}_u(\Omega_u; \Theta) = \frac{1}{|\Omega_u|} \sum_{(\mathbf{x}_i, C_i) \in \Omega_u} \mathfrak{R}_u((\mathbf{x}_i, C_i); \Theta). \quad (7)$$

Besides, inspired by that the deep feature space usually is linear and includes some meaningful semantic directions (Wang et al. 2019, 2022b), we construct semantic-level transformations based on those semantic directions, which is complementary to the pixel-level transformations, and perform semantic consistency regularization on them, so as to further regularize the classifier  $f(\cdot; \Theta)$  in the semantic level. Let the classifier  $f(\cdot; \Theta) = \mathbf{W}^\top g(\cdot; \Phi)$ ,  $g(\cdot; \Phi)$  be the deep feature extractor, and  $\mathbf{W} = [\mathbf{w}_j]_{j \in \mathcal{Y}}^\top$  be parameters of the last full-connected predictive layer. Specifically, we suppose that those semantic directions are drawn from a set of label-specific zero-mean Gaussian distributions  $\{\mathcal{N}(\mathbf{0}, \lambda \Sigma_j)\}_{j \in \mathcal{Y}}$ , where  $\lambda > 0$  controls the strength of semantic transformations. Given the known labels, we can apply those sampled label-specific semantic directions on the deep features of instances to construct semantic-level transformations. Thanks to the properties of Gaussian distribution, we can draw the semantic-transformed feature of any instance  $(\mathbf{x}_i, y_i)$  as:

$$\hat{\mathbf{a}}_i \sim \mathcal{N}(\mathbf{a}_i, \lambda \Sigma_{y_i}), \quad \mathbf{a}_i = g(\mathbf{x}_i; \Phi). \quad (8)$$

Nevertheless, the true labels of pseudo-unlabeled instances within  $\Omega_u$  are totally unknown. For each instance  $(\mathbf{x}_i, C_i) \in \Omega_u$ , we approximate its pseudo label  $\hat{y}_i$  with the CAVs on its weakly-augmented variant  $\mathbf{x}_i^w$  as:

$$\hat{y}_i = \arg \max_{j \in C_i} \hat{v}_{ij}, \quad \hat{v}_{ij} = \hat{z}_{ij} |\hat{z}_{ij} - 1|, \quad \hat{\mathbf{z}}_i = f(\mathbf{x}_i^w; \hat{\Theta}).$$

We can then draw the semantic-transformed features of its weakly- and strongly-augmented variants by applying Eq.(8). Drawing  $K$  semantic-transformed features for each augmented variant, the consistency regularization term in Eq.(6) can be rewritten as the following semantic consistency regularization term:

$$\mathfrak{R}_u^K((\mathbf{x}_i, C_i); \Theta) = \frac{1}{K^2} \sum_{k_1, k_2=1}^K h(\hat{\mathbf{p}}_i^{w, k_1}) \text{KL}(\hat{\mathbf{p}}_i^{k_1} \parallel \mathbf{p}_i^{s, k_2}),$$

s.t.  $\hat{\mathbf{a}}_i^{w, k_1} \sim \mathcal{N}(\hat{\mathbf{a}}_i^w, \lambda \Sigma_{\hat{y}_i})$ ,  $\hat{\mathbf{a}}_i^w = g(\mathbf{x}_i^w; \hat{\Phi})$ ,

$\hat{\mathbf{a}}_i^{s, k_2} \sim \mathcal{N}(\hat{\mathbf{a}}_i^s, \lambda \Sigma_{\hat{y}_i})$ ,  $\hat{\mathbf{a}}_i^s = g(\mathbf{x}_i^s; \hat{\Phi})$ ,  $(9)$

where

$$\hat{p}_{ij}^{w, k_1} = \frac{e^{\hat{z}_{ij}^{w, k_1}}}{\sum_{j' \in \mathcal{Y}} e^{\hat{z}_{ij'}^{w, k_1}}}, \quad \hat{\mathbf{z}}_i^{w, k_1} = \hat{\mathbf{W}}^\top \hat{\mathbf{a}}_i^{w, k_1};$$

$$p_{ij}^{s, k_2} = \frac{e^{z_{ij}^{s, k_2}}}{\sum_{j' \in \mathcal{Y}} e^{z_{ij'}^{s, k_2}}}, \quad \mathbf{z}_i^{s, k_2} = \mathbf{W}^\top \mathbf{a}_i^{s, k_2},$$

and further  $\hat{\mathbf{p}}_i^{k_1} = [\hat{p}_{ij}^{k_1}]_{j \in \mathcal{Y}}^\top$  is calculated as follows:

$$\hat{p}_{ij}^{k_1} = \frac{\mathbb{1}(j \in C_i) \hat{p}_{ij}^{w, k_1}}{\sum_{j' \in \mathcal{Y}} \mathbb{1}(j' \in C_i) \hat{p}_{ij'}^{w, k_1}}.$$

To avoid inefficiently sampling, we consider the expectation of Eq.(9) with all possible semantic-transformed features:

$$\mathfrak{R}_u^\infty((\mathbf{x}_i, C_i); \Theta) = \mathbb{E}_{\hat{\mathbf{a}}_i^{w, k_1}, \hat{\mathbf{a}}_i^{s, k_2}} [h(\hat{\mathbf{p}}_i^{w, k_1}) \text{KL}(\hat{\mathbf{p}}_i^{k_1} \parallel \mathbf{p}_i^{s, k_2})]. \quad (10)$$

Unfortunately, it is intractable to optimize Eq.(10) in its exact form. Alternatively, we derive an easy-to-compute upper bound  $\bar{\mathfrak{R}}_u^\infty((\mathbf{x}_i, C_i); \Theta)$  given in the following proposition. Finally, the consistency regularization over  $\Omega_u$  in Eq.(7) is rewritten below:

$$\bar{\mathcal{R}}_u(\Omega_u; \Theta) = \frac{1}{|\Omega_u|} \sum_{(\mathbf{x}_i, C_i) \in \Omega_u} \bar{\mathfrak{R}}_u^\infty((\mathbf{x}_i, C_i); \Theta). \quad (11)$$

**Proposition 1.** Suppose that  $\hat{\mathbf{a}}_i^{w, k_1} \sim \mathcal{N}(\hat{\mathbf{a}}_i^w, \lambda \Sigma_{\hat{y}_i})$  and  $\hat{\mathbf{a}}_i^{s, k_2} \sim \mathcal{N}(\hat{\mathbf{a}}_i^s, \lambda \Sigma_{\hat{y}_i})$ . Then we have an upper bound for  $\mathfrak{R}_u^\infty((\mathbf{x}_i, C_i); \Theta)$  given by

$$\mathfrak{R}_u^\infty((\mathbf{x}_i, C_i); \Theta) \leq h(\hat{\mathbf{p}}_i^w) \text{KL}(\hat{\mathbf{p}}_i \parallel \mathbf{p}_i^s) \triangleq \bar{\mathfrak{R}}_u^\infty((\mathbf{x}_i, C_i); \Theta)$$

$$\text{s.t. } \hat{\mathbf{a}}_i^w = g(\mathbf{x}_i^w; \hat{\Phi}), \quad \hat{\mathbf{a}}_i^s = g(\mathbf{x}_i^s; \hat{\Phi}),$$

where  $\hat{p}_{ij}^w = \frac{1}{-l + \sum_{j' \in \mathcal{Y}} 1/\Phi\left(\frac{\beta \hat{\mathbf{w}}_{jj'}^\top \hat{\mathbf{a}}_i^w}{(1 + \lambda \beta^2 \hat{\mathbf{w}}_{jj'}^\top \Sigma_{\hat{y}_i} \hat{\mathbf{w}}_{jj'})^{1/2}}\right)}$ ,  $\hat{p}_{ij}^s = \frac{\mathbb{1}(j \in C_i) \hat{p}_{ij}^{w, s}}{\sum_{j' \in \mathcal{Y}} \mathbb{1}(j' \in C_i) \hat{p}_{ij'}^{w, s}}$ ,  $\hat{p}_{ij}^s = \frac{e^{\mathbf{w}_j^\top \hat{\mathbf{a}}_i^s}}{\sum_{j' \in \mathcal{Y}} e^{\mathbf{w}_{j'}^\top \hat{\mathbf{a}}_i^s + \frac{1}{2} \mathbf{u}_{j'}^\top \Sigma_{\hat{y}_i} \mathbf{u}_{j'}}}$ ,  $\hat{\mathbf{u}}_{jj'} = \hat{\mathbf{w}}_j - \hat{\mathbf{w}}_{j'}$ ,  $\mathbf{u}_{j'j} = \mathbf{w}_{j'} - \mathbf{w}_j$ , and  $\Phi(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^z e^{-t^2/2} dt$  is the cumulative distribution function of the standard normal distribution  $\mathcal{N}(0, 1)$ .

## Objective of PLSP and Iterative Training Summary

We summarize the overall objective of PLSP, and clarify the training details in following.

**Objective of PLSP.** Besides the aforementioned pseudo-supervised loss  $\mathcal{L}_l$  and regularization  $\mathcal{R}_u$ , we also incorporate a complementary loss over  $\Omega_u$  to minimize the predictions of non-candidate labels:

$$\mathcal{L}_{cl}(\Omega_u; \Theta) = \frac{1}{|\Omega_u|} \sum_{(\mathbf{x}_i, C_i) \in \Omega_u} \sum_{j \notin C_i} -\log(1 - p_{ij}). \quad (12)$$

And we also improve  $\mathcal{L}_l$  and  $\mathcal{L}_{cl}$  with the semantic-level transformation Eq.(8), then obtain their corresponding upper bounds according to Proposition 1, given by:

$$\bar{\mathcal{L}}_l(\Omega_l; \Theta) = \frac{1}{|\Omega_l|} \sum_{(\mathbf{x}_i, y_i) \in \Omega_l} -\log p_{iy_i}, \quad (13)$$

$$\bar{\mathcal{L}}_{cl}(\Omega_u; \Theta) = \frac{1}{|\Omega_u|} \sum_{(\mathbf{x}_i, C_i) \in \Omega_u} \sum_{j \notin C_i} -\log(1 - p_{ij}), \quad (14)$$

where

$$p_{iy_i} = \frac{e^{\mathbf{w}_{y_i}^\top \mathbf{a}_i}}{\sum_{j' \in \mathcal{Y}} e^{\mathbf{w}_{j'}^\top \mathbf{a}_i + \frac{\lambda}{2} \mathbf{u}_{j'}^\top \Sigma_{y_i} \mathbf{u}_{j'} y_i}}, \quad \mathbf{a}_i = g(\mathbf{x}_i, \Phi),$$

$$p_{ij} = \frac{e^{\mathbf{w}_j^\top \mathbf{a}_i}}{\sum_{j' \in \mathcal{Y}} e^{\mathbf{w}_{j'}^\top \mathbf{a}_i + \frac{\lambda}{2} \mathbf{u}_{j'}^\top \Sigma_{\hat{y}_i} \mathbf{u}_{j'} j}}, \quad \mathbf{a}_i = g(\mathbf{x}_i, \Phi),$$

and  $\mathbf{u}_{j' y_i} = \mathbf{w}_{j'} - \mathbf{w}_{y_i}$ . Accordingly, the final objective of PLSP can be reformulated as:

$$\mathcal{L}(\{\Omega_l, \Omega_u\}; \Theta) = \gamma (\bar{\mathcal{L}}_l(\Omega_l; \Theta) + \bar{\mathcal{R}}_u(\Omega_u; \Theta)) + \bar{\mathcal{L}}_{cl}(\Omega_u; \Theta), \quad (15)$$

where  $\gamma > 0$  is the hyper-parameter to balance the SS learning loss and complementary loss.

**Update of label-specific covariance matrices**  $\{\Sigma_j\}_{j \in \mathcal{Y}}$ . Following (Wang et al. 2019, 2022b), we approximate  $\{\Sigma_j\}_{j \in \mathcal{Y}}$  with pseudo-labeled instances by counting statistics from all mini-batches incrementally. For each  $\Sigma_j$  in the  $c$ -th iteration, it can be updated as follows:

$$\begin{aligned} \Sigma_j^{(c)} &= \frac{m_j^{(c-1)} \Sigma_j^{(c-1)} + m_j'^{(c)} \Sigma_j'^{(c)}}{m_j^{(c-1)} + m_j'^{(c)}} \\ &+ \frac{m_j^{(c-1)} m_j'^{(c)} (\boldsymbol{\mu}_j^{(c-1)} - \boldsymbol{\mu}_j'^{(c)}) (\boldsymbol{\mu}_j^{(c-1)} - \boldsymbol{\mu}_j'^{(c)})^\top}{(m_j^{(c-1)} + m_j'^{(c)})^2}, \quad (16) \\ \boldsymbol{\mu}_j^{(c)} &= \frac{m_j^{(c-1)} \boldsymbol{\mu}_j^{(c-1)} + m_j'^{(c)} \boldsymbol{\mu}_j'^{(c)}}{m_j^{(c-1)} + m_j'^{(c)}}, \quad m_j^{(c)} = m_j^{(c-1)} + m_j'^{(c)}, \end{aligned}$$

where  $\boldsymbol{\mu}_j'^{(c)}$  and  $\Sigma_j'^{(c)}$  are the mean and covariance matrix of features within class  $j$  in  $c$ -th mini-batch, respectively;  $m_j^{(c)}$  the total number of pseudo-labeled instances belonging to class  $j$  in all  $c$  mini-batches and  $m_j'^{(c)}$  the number of pseudo-labeled instances belonging to class  $j$  in  $c$ -th mini-batch.

**Adjusting the SS learning loss weight**  $\gamma$ . In the early training stage, the SS learning loss may be less accurate. To fix issue, We dynamically adjust the SS learning loss weight  $\gamma$  by a non-decreasing function  $\gamma = \min\{\frac{t}{T} \gamma_0, \gamma_0\}$  with respect to the epoch number  $t$ , where  $\gamma_0$  is the maximum weight, and  $T$  the maximum number of SS training epochs.

**Adjusting the confidence threshold**  $\tau$ . We employ the curriculum pseudo labeling (Zhang et al. 2021) to adjust  $\tau$ . For each class  $j$ , its value at  $c$ -th iteration is calculated by:

$$\begin{aligned} \tau_c(j) &= \eta_c(j) \cdot \tau_0, \quad \eta_c(j) = \frac{\sigma_c(j)}{\max_{j' \in \mathcal{Y}} \sigma_c(j')}, \\ \sigma_c(j) &= \sum_{(\mathbf{x}_i, C_i) \in \Omega_u} h(\hat{\mathbf{p}}_i^w) \cdot \mathbf{1}(\hat{y}_i = j), \end{aligned}$$

where  $\tau_0$  is the maximum confidence threshold.

**Adjusting the transformation strength**  $\lambda$ . Following (Wang et al. 2019, 2022b), we dynamically adjust the transformation strength  $\lambda$  with a non-decreasing function  $\lambda = \min\{\frac{t}{T} \lambda_0, \lambda_0\}$  with respect to the epoch number  $t$ , where  $\lambda_0$  is the maximum transformation strength, so as to reduce the negative impact of the low-quality estimations of covariance matrices in the early training stage.

Algorithm 1: Training procedure of PLSP

**Input:**  $\Omega$ : PL training dataset  $\Omega = \{(\mathbf{x}_i, C_i)\}_{i=1}^n$ ;  $m$ : number of pseudo-labeled instances;  $\gamma_0$ : SS learning loss weight;  $\tau_0$ : confidence threshold;  $\lambda_0$ : semantic transformation strength; **Output:**  $\Theta$ : classifier parameters

```

1: Initialize the classifier parameters  $\Theta = \{\Phi, \mathbf{W}\}$ ;
2: for  $t = 0$  to  $T_0$  do {% Pre-training stage %}
3:   for  $c = 0$  to  $I$  do
4:     Sample a mini-batch  $\{(\mathbf{x}_i, C_i)\}_{i=1}^B$  from  $\Omega$ ;
5:     Compute  $\mathcal{L}_{df}$  according to Eq.(2);
6:     Update  $\Theta$  with SGD;
7:   end for
8: end for
9: for  $t = 0$  to  $T$  do {% SS training stage %}
10:  Construct pseudo-training dataset  $\{\Omega_l, \Omega_u\}$  according to Eqs.(3) and (4);
11:  for  $c = 0$  to  $I$  do
12:    Sample a mini-batch  $\{(\mathbf{x}_i, y_i)\}_{i=1}^{B_l}$  from  $\Omega_l$  and a mini-batch  $\{(\mathbf{x}_i^w, \mathbf{x}_i^s, C_i)\}_{i=1}^{B_u}$  from  $\Omega_u$  with  $\alpha(\cdot)$  and  $\mathcal{A}(\cdot)$ ;
13:    Compute  $\mathbf{a}_i = g(\mathbf{x}_i; \Phi)$ ,  $\mathbf{a}_i^w = g(\mathbf{x}_i^w; \Phi)$ ,  $\mathbf{a}_i^s = g(\mathbf{x}_i^s; \Phi)$ ;
14:    Estimate covariance matrices  $\{\Sigma_j\}_{j \in \mathcal{Y}}$  according to Eq.(16);
15:    Compute  $\mathcal{L}$  according to Eq.(15);
16:    Update  $\Theta$  with SGD;
17:  end for
18: end for

```

**Iterative training summary.** In practice, to prevent the error memorization and reduce the time cost, we update  $\{\Omega_l, \Omega_u\}$  with the current predictions per-epoch. The classifier parameters  $\Theta$  are optimized by using the stochastic optimization with SGD. Overall, the iterative training procedure of PLSP is summarized in *Algorithm 1*.

## Experiment

### Experimental Setup

**Datasets.** We utilize 3 widely used benchmark image datasets, including Fashion-MNIST (Xiao, Rasul, and Vollgraf 2017), CIFAR-10 and CIFAR-100 (Krizhevsky 2016). We manually synthesize the partially labeled versions of these datasets by applying Uniformly Sampling Strategy (USS) (Feng et al. 2020) and Flipping Probability Strategy (FPS) (Lv et al. 2020). The former one is conducted by uniformly sampling a candidate label set from the candidate label set space  $\mathcal{C}$  for each instance, and the latter one generates the candidate label set of each instance by selecting any irrelevant label as its candidate one with a flipping probability  $q$ .<sup>4</sup> In experiments, we employ  $q \in \{0.3, 0.5, 0.7\}$  for Fashion-MNIST and CIFAR-10, and  $q \in \{0.05, 0.1, 0.2\}$  for CIFAR-100 due to the more labels. We adopt 5-layer LeNet, 22-layer Densenet and 18-layer ResNet as the backbones of Fashion-MNIST, CIFAR-10 and CIFAR-100, respectively.

<sup>4</sup>Note that the flipping probability strategy will uniformly flip a random irrelevant label into the candidate label set when none of irrelevant labels are flipped.

Metric	Dataset	CC	RC	PRODEN	LWS	CAVL	PiCO	PLSP
(I) USS								
Macro-F1	FMNIST	0.879±0.001 <sup>‡</sup>	0.893±0.001	0.891±0.006	0.881±0.007 <sup>‡</sup>	0.882±0.004 <sup>‡</sup>	<b>0.907±0.001</b>	0.897±0.003
	CIFAR-10	0.745±0.006 <sup>‡</sup>	0.787±0.003 <sup>‡</sup>	0.796±0.003 <sup>‡</sup>	0.781±0.007 <sup>‡</sup>	0.748±0.051 <sup>‡</sup>	0.869±0.001 <sup>‡</sup>	<b>0.889±0.003</b>
Micro-F1	FMNIST	0.880±0.001 <sup>‡</sup>	0.893±0.001	0.892±0.003	0.883±0.005 <sup>‡</sup>	0.883±0.004 <sup>‡</sup>	<b>0.907±0.001</b>	0.897±0.002
	CIFAR-10	0.747±0.005 <sup>‡</sup>	0.788±0.003 <sup>‡</sup>	0.796±0.003 <sup>‡</sup>	0.782±0.004 <sup>‡</sup>	0.756±0.038 <sup>‡</sup>	0.870±0.001 <sup>‡</sup>	<b>0.889±0.003</b>
(II) FPS ( $q = 0.3$ )								
Macro-F1	FMNIST	0.883±0.003 <sup>‡</sup>	0.893±0.001	0.894±0.003	0.888±0.005	0.887±0.003	<b>0.909±0.003</b>	0.894±0.002
	CIFAR-10	0.769±0.005 <sup>‡</sup>	0.803±0.003 <sup>‡</sup>	0.801±0.005 <sup>‡</sup>	0.802±0.005 <sup>‡</sup>	0.796±0.004 <sup>‡</sup>	0.880±0.003 <sup>‡</sup>	<b>0.898±0.002</b>
Micro-F1	FMNIST	0.883±0.003	0.894±0.001	0.895±0.002	0.889±0.004	0.887±0.003	<b>0.909±0.003</b>	0.894±0.002
	CIFAR-10	0.769±0.004 <sup>‡</sup>	0.803±0.003 <sup>‡</sup>	0.801±0.005 <sup>‡</sup>	0.803±0.004 <sup>‡</sup>	0.796±0.004 <sup>‡</sup>	0.880±0.003 <sup>‡</sup>	<b>0.898±0.002</b>
(III) FPS ( $q = 0.5$ )								
Macro-F1	FMNIST	0.881±0.001	0.890±0.002	0.891±0.004	0.884±0.006	0.881±0.003	<b>0.903±0.002</b>	0.891±0.003
	CIFAR-10	0.734±0.007 <sup>‡</sup>	0.782±0.004 <sup>‡</sup>	0.791±0.004 <sup>‡</sup>	0.794±0.003 <sup>‡</sup>	0.767±0.004 <sup>‡</sup>	0.865±0.002 <sup>‡</sup>	<b>0.887±0.002</b>
Micro-F1	FMNIST	0.881±0.001 <sup>‡</sup>	0.891±0.002	0.892±0.002	0.885±0.004	0.882±0.004	<b>0.903±0.002</b>	0.892±0.002
	CIFAR-10	0.735±0.008 <sup>‡</sup>	0.783±0.004 <sup>‡</sup>	0.791±0.004 <sup>‡</sup>	0.795±0.002 <sup>‡</sup>	0.768±0.004 <sup>‡</sup>	0.866±0.002 <sup>‡</sup>	<b>0.887±0.002</b>
(IV) FPS ( $q = 0.7$ )								
Macro-F1	FMNIST	0.874±0.007	<b>0.886±0.006</b>	0.884±0.003	0.875±0.002	0.863±0.002 <sup>‡</sup>	0.865±0.035 <sup>‡</sup>	0.878±0.005
	CIFAR-10	0.678±0.008 <sup>‡</sup>	0.728±0.002 <sup>‡</sup>	0.745±0.004 <sup>‡</sup>	0.743±0.005 <sup>‡</sup>	0.673±0.051 <sup>‡</sup>	0.808±0.045 <sup>‡</sup>	<b>0.869±0.006</b>
Micro-F1	FMNIST	0.875±0.007	<b>0.886±0.002</b>	0.885±0.001	0.875±0.001	0.853±0.002 <sup>‡</sup>	0.872±0.020	0.879±0.004
	CIFAR-10	0.681±0.008 <sup>‡</sup>	0.730±0.002 <sup>‡</sup>	0.746±0.004 <sup>‡</sup>	0.744±0.005 <sup>‡</sup>	0.692±0.035 <sup>‡</sup>	0.816±0.032 <sup>‡</sup>	<b>0.870±0.006</b>

Table 1: Empirical results (mean±std) on Fashion-MNIST (FMNIST) and CIFAR-10 with different data generation strategies and ambiguity levels: (I) USS; (II) FPS ( $q=0.3$ ); (III) FPS ( $q=0.5$ ); (IV) FPS ( $q=0.7$ ). The highest scores are indicated in bold. The notation “<sup>‡</sup>” indicates that the performance gain of PLSP is statistically significant (paired sample t-tests) at 0.01 level.

**Baseline PL learning methods and training settings.** We compare PLSP against the following 6 existing deep PL learning methods, including RC (Feng et al. 2020), CC (Feng et al. 2020), PRODEN (Lv et al. 2020), LW (Wen et al. 2021) with sigmoid loss function, CAVL (Zhang et al. 2022), and PiCO (Wang et al. 2022a). We train all methods by using the SGD optimizer, and search the learning rate from  $\{0.0001, 0.001, 0.01, 0.05, 0.1, 0.5\}$  and the weight decay from  $\{10^{-6}, 10^{-5}, \dots, 10^{-1}\}$ . For all baselines and the pre-training-stage of PLSP, we set the batch size 256 for Fashion-MNIST and CIFAR-10, and 64 for CIFAR-100. For all baselines, we employ the default or suggested settings of hyper-parameters in their papers and released codes. For PLSP, we use the following hyper-parameter settings:  $\gamma_0 = 1.0$ ,  $\lambda_0 = 0.01$ ,  $\tau_0 = 0.75$ , number of pre-training epochs  $T_0 = 10$ , number of SS training epochs  $T = 250$ , number of inner loops  $I = 200$ , batch sizes of pseudo-labeled and pseudo-unlabeled instances  $B_l = 64$ ,  $B_u = 256$ . Specially, for CIFAR-100 we set  $T_0 = 50$ ,  $I = 800$ ,  $B_l = 16$ ,  $B_u = 64$ . We set the number of pseudo-labeled instances per-class  $k = 200$ . Besides, we employ the horizontal flipping and cropping to conduct the weakly augmentation function  $\alpha(\cdot)$  of all datasets, and implement the strongly augmentation function  $\mathcal{A}(\cdot)$  for Fashion-MNIST with horizontal flipping, cropping and Cutout, for CIFAR-10 and CIFAR-100 with horizontal flipping, cropping, Cutout as well as AutoAugment.<sup>5</sup> All experiments are carried on a Linux server with one NVIDIA GeForce RTX 3090 GPU.

<sup>5</sup>For AutoAugment, we simply utilize the augmentation policies released by (Cubuk et al. 2019).

**Evaluation metrics.** We employ Macro-F1 and Micro-F1 to evaluate the classification performance, and calculate them by using the Scikit-Learn tools (Pedregosa et al. 2011).

## Main Results

We perform all experiments with five different random seeds, and report the average scores of Fashion-MNIST and CIFAR-10 in Table 1, and ones of CIFAR-100 in Table 2. Overall, our PLSP significantly outperforms all comparing methods in most cases, and achieves particularly significant performance gain on high ambiguity levels. As shown in Tables 1 and 2: (1) Our PLSP consistently perform better than all baselines on CIFAR-10 and achieves a competitive performance on Fashion-MNIST across four partial label settings. For example, Micro-F1 scores of PLSP are 0.019 ~ 0.054 higher than ones of the recent state-of-the-art PiCO on four partially-labeled versions of CIFAR-10, and even gain 0.054 significant improvement on high ambiguity level, *i.e.*,  $q = 0.7$ . (2) Compared with all baselines, our PLSP achieves very significant performance gain on CIFAR-100 across  $q = 0.05, 0.1$  and 0.2, and show more significant superiority than that on previous simpler datasets. (3) Besides, PiCO always drop dramatically on Fashion-MNIST and CIFAR-10 with  $q = 0.7$ , especially CIFAR-100 with  $q = 0.2$ . The possible reason is that PiCO could not identify true labels with contrastive representation learning to disambiguate candidate labels when on high ambiguity level.

## Ablation Study

In this section, we perform extensive experiments to examine the importance of different components of PLSP. We compare

Metric	$q$	CC	RC	PRODEN	LWS	CAVL	PiCO	PLSP
Macro-F1	0.05	0.469±0.003 <sup>‡</sup>	0.461±0.007 <sup>‡</sup>	0.601±0.004 <sup>‡</sup>	0.567±0.008 <sup>‡</sup>	0.398±0.008 <sup>‡</sup>	0.744±0.007 <sup>‡</sup>	<b>0.770±0.002</b>
	0.1	0.431±0.006 <sup>‡</sup>	0.388±0.006 <sup>‡</sup>	0.512±0.006 <sup>‡</sup>	0.498±0.005 <sup>‡</sup>	0.229±0.018 <sup>‡</sup>	0.636±0.021 <sup>‡</sup>	<b>0.733±0.013</b>
	0.2	0.348±0.008 <sup>‡</sup>	0.230±0.013 <sup>‡</sup>	0.476±0.010 <sup>‡</sup>	0.401±0.015 <sup>‡</sup>	0.066±0.010 <sup>‡</sup>	0.190±0.025 <sup>‡</sup>	<b>0.660±0.008</b>
Micro-F1	0.05	0.470±0.004 <sup>‡</sup>	0.465±0.006 <sup>‡</sup>	0.607±0.001 <sup>‡</sup>	0.596±0.003 <sup>‡</sup>	0.402±0.008 <sup>‡</sup>	0.746±0.006 <sup>‡</sup>	<b>0.770±0.002</b>
	0.1	0.435±0.005 <sup>‡</sup>	0.400±0.005 <sup>‡</sup>	0.568±0.003 <sup>‡</sup>	0.535±0.001 <sup>‡</sup>	0.262±0.015 <sup>‡</sup>	0.660±0.015 <sup>‡</sup>	<b>0.739±0.008</b>
	0.2	0.357±0.009 <sup>‡</sup>	0.279±0.012 <sup>‡</sup>	0.496±0.007 <sup>‡</sup>	0.434±0.011 <sup>‡</sup>	0.104±0.009 <sup>‡</sup>	0.288±0.022 <sup>‡</sup>	<b>0.687±0.006</b>

Table 2: Empirical results (mean±std) on CIFAR-100 with FPS ( $q = 0.05, 0.1, 0.2$ ). The highest scores are indicated in bold. The notation “<sup>‡</sup>” indicates that the performance gain of PLSP is statistically significant (paired sample t-tests) at 0.01 level.

Method	CIFAR-10	
	Macro-F1	Micro-F1
PLSP	<b>0.869±0.006</b>	<b>0.870±0.006</b>
PLSP w/o ST	0.858±0.008	0.861±0.007
DF	0.413±0.015	0.418±0.012

Table 3: Ablation study results (mean±std) on CIFAR-10 with FPS ( $q = 0.7$ ). The highest scores are indicated in bold.

PLSP with PLSP without the semantic transformation (ST) and the version training only with the disambiguation-free (DF) objective of Eq.(2) on CIFAR-10 by using data generation with FPS on  $q = 0.7$ . The experimental results are reported in Table 3. It clearly demonstrates that the proposed SS learning strategy can significantly improve the classification performance of PL learning. Besides, we can also observe that the semantic transformation can also improve the classification performance, proving its effectiveness to capture the semantic consistency.

### Sensitivity Analysis

In this section, we examine the sensitivities of number of pseudo-instances per-class  $k$ . We conduct the sensitive experiments by varying  $k$  over  $\{0, 50, 100, 200, 500, 1000, 5000\}$  on CIFAR-10 by using data generation with FPS on  $q = 0.7$ , and illustrate the experimental results in Fig.2. As is shown: (1) Obviously, the performance is relatively stable when  $k \leq 200$  and achieve the highest when  $k = 200$ , and it sharply drops as the values become bigger. It is expected since the smaller value of  $k$  may ignore some high-confidence instances and the bigger value of  $k$  will introduce many “unreliable” pseudo-labeled instances, leading to a poor classifier. (2) Moreover, the performance is poor when both  $k = 0$  and 5000, especially when  $k = 5000$ . Notice that when  $k = 0$  none of instances within  $\Omega$  are selected as pseudo-labeled instances, *i.e.*,  $\{\Omega_l = \emptyset, \Omega_u = \Omega\}$ , and when  $k = 5000$  for CIFAR-10 all instances are selected as pseudo-labeled instances, *i.e.*,  $\{\Omega_l = \Omega, \Omega_u = \Omega\}$ . It demonstrates the effectiveness of the proposed SS learning strategy for PL learning task. In practice, we suggest tuning  $k$  over the set  $\{50, 100, 200\}$ .

### Efficiency Comparison

To examine the efficiency of our PLSP, we perform efficiency comparisons over PLSP and PiCO on all benchmarks with USS. We compare the overall time costs during pre-training and training stages respectively, and perform experiments

Method	PiCO		PLSP	
	Pretrain	Train	Pretrain	Train
FMNIST	–	24,400s	60s	10,920s
CIFAR-10	–	38,000s	107s	25,000s
CIFAR-100	–	55,200s	299s	44,000s

Table 4: Time cost (second, s) of PLSP and PiCO on Fashion-MNIST (FMNIST), CIFAR-10 and CIFAR-100 with USS.




Figure 2: Sensitivity analysis of the number of pseudo-labeled instances per-class  $k$  on CIFAR-10 with FPS ( $q = 0.7$ ).

with the suggested settings for all methods and benchmarks. Table 4 shows the running time results averaged on 10 runs. As is shown: (1) Obviously, the additional disambiguation-free pre-training stage of PLSP is very efficient. (2) Moreover, In contrast to PiCO, PLSP empirically converges fast due to more reliable supervision with SSL perspective (PLSP 250 epochs vs PiCO 800 epochs) and costs less time in practice during the training stage.

## Conclusion

In this work, we develop a novel PL learning method named PLSP by resolving the PL learning problem from the semi-supervised perspective with strong SS learning techniques. We conduct the SS learning strategy by selecting high-confidence partially-labeled instances as pseudo-labeled instances and treating the remained ones as pseudo-unlabeled. We design a semantic consistency regularization with respect to the semantic-transformed weakly- and strongly-augmented instances, and derive its approximation form for efficient optimization. Empirical results demonstrate the superior performance of PLSP compared with the existing PL learning baselines, especially on high ambiguity levels.

## Acknowledgments

We would like to acknowledge support for this project from the National Key R&D Program of China (No.2021ZD0112501, No.2021ZD0112502), the National Natural Science Foundation of China (NSFC) (No.62276113, No.62006094, No.61876071), the Key R&D Projects of Science and Technology Department of Jilin Province of China (No.20180201003SF, No.20190701031GH).

## References

- Berthelot, D.; Carlini, N.; Cubuk, E. D.; Kurakin, A.; Sohn, K.; Zhang, H.; and Raffel, C. 2020. ReMixMatch: Semi-Supervised Learning with Distribution Matching and Augmentation Anchoring. In *ICLR*.
- Berthelot, D.; Carlini, N.; Goodfellow, I. J.; Papernot, N.; Oliver, A.; and Raffel, C. 2019. MixMatch: A Holistic Approach to Semi-Supervised Learning. In *NeurIPS*, 5050–5060.
- Chen, C.; Patel, V. M.; and Chellappa, R. 2018. Learning from Ambiguously Labeled Face Images. *IEEE TPAMI*, 40(7): 1653–1667.
- Cour, T.; Sapp, B.; Jordan, C.; and Taskar, B. 2009. Learning from Ambiguously Labeled Images. In *IEEE CVPR*, 919–926.
- Cour, T.; Sapp, B.; and Taskar, B. 2011. Learning from Partial Labels. *JMLR*, 12(5): 1501–1536.
- Cubuk, E. D.; Zoph, B.; Mané, D.; Vasudevan, V.; and Le, Q. V. 2019. AutoAugment: Learning Augmentation Strategies From Data. In *IEEE CVPR*, 113–123.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. 2020. RandAugment: Practical Automated Data Augmentation with a Reduced Search Space. In *NeurIPS*, 18613–18624.
- 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. 2018. Leveraging Latent Label Distributions for Partial Label Learning. In *IJCAI*, 2107–2113.
- Feng, L.; and An, B. 2019a. Partial Label Learning by Semantic Difference Maximization. In *IJCAI*, 2294–2300.
- Feng, L.; and An, B. 2019b. Partial Label Learning with Self-Guided Retraining. In *AAAI*, 3542–3549.
- Feng, L.; Lv, J.; Han, B.; Xu, M.; Niu, G.; Geng, X.; An, B.; and Sugiyama, M. 2020. Provably Consistent Partial-Label Learning. In *NeurIPS*, 10948–10960.
- Krizhevsky, A. 2016. *Learning Multiple Layers of Features from Tiny Images*. Technical report, University of Toronto.
- Laine, S.; and Aila, T. 2017. Temporal Ensembling for Semi-Supervised Learning. In *ICLR*.
- Li, C.; Li, X.; Feng, L.; and Ouyang, J. 2022. Who Is Your Right Mixup Partner in Positive and Unlabeled Learning. In *ICLR*.
- Li, C.; Li, X.; and Ouyang, J. 2020. Learning with Noisy Partial Labels by Simultaneously Leveraging Global and Local Consistencies. In *ACM CIKM*, 725–734.
- Li, C.; Li, X.; and Ouyang, J. 2021. Semi-Supervised Text Classification with Balanced Deep Representation Distributions. In *ACL-IJCNLP*, 5044–5053.
- Li, C.; Li, X.; Ouyang, J.; and Wang, Y. 2021. Detecting the Fake Candidate Instances: Ambiguous Label Learning with Generative Adversarial Networks. In *ACM CIKM*, 903–912.
- Li, J.; Socher, R.; and Hoi, S. C. H. 2020. DivideMix: Learning with Noisy Labels as Semi-supervised Learning. In *ICLR*.
- Luo, J.; and Orabona, F. 2010. Learning from Candidate Labeling Sets. In *NeurIPS*, 1504–1512.
- 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*, 6500–6510.
- Miyato, T.; Maeda, S.; Koyama, M.; and Ishii, S. 2019. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning. *IEEE TPAMI*, 41(8): 1979–1993.
- Ni, P.; Zhao, S.; Dai, Z.; Chen, H.; and Li, C. 2021. Partial Label Learning via Conditional-Label-Aware Disambiguation. *JCST*, 36(3): 590–605.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; VanderPlas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *JMLR*, 12: 2825–2830.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C.; Cubuk, E. D.; Kurakin, A.; and Li, C. 2020. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. In *NeurIPS*, 596–608.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 1195–1204.
- Wang, H.; Xiao, R.; Li, Y.; Feng, L.; Niu, G.; and Zhao, J. 2022a. PiCO: Contrastive Label Disambiguation for Partial Label Learning. In *ICLR*.
- Wang, Y.; Huang, G.; Song, S.; Pan, X.; Xia, Y.; and Wu, C. 2022b. Regularizing Deep Networks With Semantic Data Augmentation. *IEEE TPAMI*, 44(7): 3733–3748.
- Wang, Y.; Pan, X.; Song, S.; Zhang, H.; Huang, G.; and Wu, C. 2019. Implicit Semantic Data Augmentation for Deep Networks. In *NeurIPS*, 12614–12623.
- Wen, H.; Cui, J.; Hang, H.; Liu, J.; Wang, Y.; and Lin, Z. 2021. Leveraged Weighted Loss for Partial Label Learning. In *ICML*, 11091–11100.
- Wu, D.; Wang, D.; and Zhang, M. 2022. Revisiting Consistency Regularization for Deep Partial Label Learning. In *ICML*, 24212–24225.
- Wu, X.; and Zhang, M. 2018. Towards Enabling Binary Decomposition for Partial Label Learning. In *IJCAI*, 2868–2874.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv preprint arXiv:1708.07747*.



- Xie, Q.; Dai, Z.; Hovy, E. H.; Luong, T.; and Le, Q. 2020. Unsupervised Data Augmentation for Consistency Training. In *NeurIPS*, 6256–6268.
- Xu, N.; Lv, J.; and Geng, X. 2019. Partial Label Learning via Label Enhancement. In *AAAI*, 5557–5564.
- Xu, N.; Qiao, C.; Geng, X.; and Zhang, M. 2021. Instance-Dependent Partial Label Learning. In *NeurIPS*, 27119–27130.
- Yan, Y.; and Guo, Y. 2020. Multi-Level Generative Models for Partial Label Learning with Non-random Label Noise. *arXiv preprint arXiv:2005.05407*.
- Zeng, Z.; Xiao, S.; Jia, K.; Chan, T.; Gao, S.; Xu, D.; and Ma, Y. 2013. Learning by Associating Ambiguously Labeled Images. In *IEEE CVPR*, 708–715.
- 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. In *NeurIPS*, 18408–18419.
- Zhang, F.; Feng, L.; Han, B.; Liu, T.; Niu, G.; Qin, T.; and Sugiyama, M. 2022. Exploiting Class Activation Value for Partial-Label Learning. In *ICLR*.
- Zhang, M.; and Yu, F. 2015. Solving the Partial Label Learning Problem: An Instance-Based Approach. In *IJCAI*, 4048–4054.
- Zhang, M.; Yu, F.; and Tang, C. 2017. Disambiguation-Free Partial Label Learning. *IEEE TKDE*, 29(10): 2155–2167.
- Zhang, M.; Zhou, B.; and Liu, X. 2016. Partial Label Learning via Feature-Aware Disambiguation. In *ACM SIGKDD*, 1335–1344.
- Zhang, M.; Zhou, B.; and Liu, X. 2019. Adaptive Graph Guided Disambiguation for Partial Label Learning. In *ACM SIGKDD*, 83–91.
- Zhang, Y.; Yang, G.; Zhao, S.; Ni, P.; Lian, H.; Chen, H.; and Li, C. 2020. Partial Label Learning via Generative Adversarial Nets. In *ECAI*, 1674–1681.