

## Two-stage Training for Learning from Label Proportions

Jiabin Liu<sup>1,4\*</sup>, Bo Wang<sup>2\*†</sup>, Xin Shen<sup>3</sup>, Zhiquan Qi<sup>4†</sup> and Yingjie Tian<sup>4</sup>

<sup>1</sup>AI Lab, Samsung Research China - Beijing

<sup>2</sup>University of International Business and Economics

<sup>3</sup>The Chinese University of Hong Kong

<sup>4</sup>University of Chinese Academy of Sciences

liujiabin008@126.com, wangbo@uibe.edu.cn, xshen@se.cuhk.edu.hk,  
qizhiquan@foxmail.com, tyj@ucas.ac.cn

### Abstract

Learning from label proportions (LLP) aims at learning an instance-level classifier with label proportions in grouped training data. Existing deep learning based LLP methods utilize end-to-end pipelines to obtain the proportional loss with Kullback-Leibler divergence between the bag-level prior and posterior class distributions. However, the unconstrained optimization on this objective can hardly reach a solution in accordance with the given proportions. Besides, concerning the probabilistic classifier, this strategy unavoidably results in high-entropy conditional class distributions at the instance level. These issues further degrade the performance of the instance-level classification. In this paper, we regard these problems as noisy pseudo labeling, and instead impose the strict proportion consistency on the classifier with a constrained optimization as a continuous training stage for existing LLP classifiers. In addition, we introduce the mixup strategy and symmetric cross-entropy to further reduce the label noise. Our framework is model-agnostic, and demonstrates compelling performance improvement in extensive experiments, when incorporated into other deep LLP models as a post-hoc phase.

### 1 Introduction

Learning from label proportions (LLP) is an important weakly supervised classification problem with only the label proportions in grouped data available. Still, training LLP aims to obtain an instance-level classifier for the new-come inputs. Successfully resolving LLP problems greatly contribute to many real-life applications: demographic classification [Ardehaly and Culotta, 2017], US presidential election [Sun *et al.*, 2017], embryo implantation prediction [Hernández-González *et al.*, 2018], spam filtering [Kuck and de Freitas, 2012], video event detection [Lai *et al.*, 2014], visual attribute modeling [Chen *et al.*, 2014; Yu *et al.*, 2014a], and traffic flow prediction [Liebig *et al.*, 2015].

\*Equal Contribution

†Corresponding Authors

On the one hand, the learnability of LLP strongly depends on the instances grouping and the proportions distribution in the bags. [Yu *et al.*, 2014b] studied the instance-level empirical proportions risk minimization (EPRM) algorithm for LLP, w.r.t. the number, the size, and the prior label distribution of the bags. They proved that LLP is learnable with the EPRM principle and given the bound of expected learning risk.

On the other hand, EPRM strives to minimize bag-level label proportions error. Normally, this goal is achieved by minimizing Kullback-Leibler (KL) divergence between prior and posterior class distributions in each bag. However, bag-level proportional information hardly provides sufficient constraints to perfectly solve LLP, because too many instance-level classifiers can satisfy proportional constraints exactly. In other words, when considering instance-level classification, LLP is ill-posed. As a consequence of the underdetermination, despite a number of achievements have been developed to resolve LLP accurately, it is still of great importance to design effective instance-level learning scheme to significantly improve the performance on high-dimensional data, e.g., images, merely with the proportional information.

In this paper, we challenge this issue through explicit *weakly supervised clustering* with proportions, instead of the inadequate bag-level KL divergence. To be concrete, unsupervised clustering is expected to discover data clusters specifically corresponding to classes. However, naively applying clustering without supervision will result in a trivial solution to assign all the instances to one cluster, deviating from the semantics, i.e., classes. Fortunately, we can avoid this degeneration by imposing appropriate constraints on the cluster distribution [Caron *et al.*, 2018]. For example, when there is no knowledge on labels, we impose discrete uniform distribution to labels, leading to equal clustering. Besides, in a semi-supervised learning protocol, the clustering result can be restricted with the help of labeled instances [Asano *et al.*, 2020] as well. Similarly, in LLP scenario, we draw inspiration from this constrained clustering, to leverage label proportions to constrain cluster sizes. Specifically, we build a constrained optimization problem, regarding feasible solutions as pseudo-labels that accurately complying proportions.

In order to properly tackle the weak supervision, we cast our framework, called PLOT, by considering pseudo labeling and the classification within one objective, utilizing label proportions as the constraints. To implement the proposed

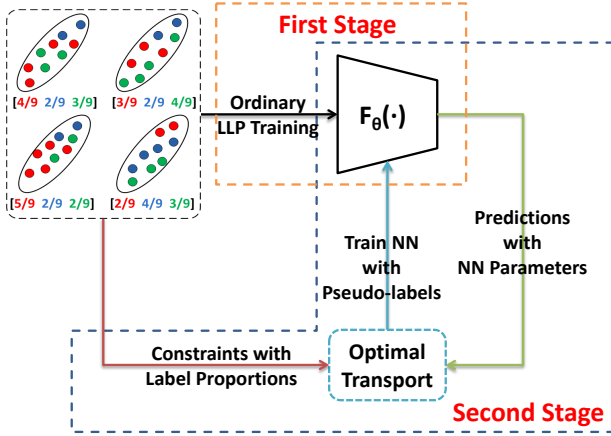


Figure 1: Illustration of the proposed two-stage LLP training framework. (Best viewed in color)

schema, we alternately update the network parameters and the pseudo-labels, where optimal transport (OT) [Peyré *et al.*, 2019] is employed to conduct the pseudo labeling process and standard cross-entropy minimization is adopted to train the network parameters. Besides, to recognize the resulting pseudo-label as noisy label to the unknown clean one, we spur a new line to propose a *two-stage* training strategy: Set our method as a following stage after other LLP solvers, and train the classification network with *symmetric cross-entropy (SCE)* [Wang *et al.*, 2019] along with *mixup* dataset [Zhang *et al.*, 2018] to combat the memorization of noisy labels and inject the robustness into the solution. We illustrate the proposed two-stage framework in Figure 1.

In summary, our main contributions are four-fold:

- We propose a novel LLP framework PLOT, which aims at meeting the exact proportions with a constrained optimization, and is an essentially orthogonal LLP treatment to the existing LLP algorithms;
- We conduct an alternate optimization process to solve the optimization on both the neural network classifier and pseudo labeling, by applying SCE to learn from the noisy pseudo-labels obtained with OT;
- Our framework is model-agnostic, and we demonstrate that it can easily fit for various deep-based LLP models as a post-hoc stage to further boost their performance;
- With lightly additional hyper-parameter and mixup data augmentation, our framework achieves state-of-the-art LLP performance on several benchmark datasets, based on the neural network pipeline with standard settings.

## 2 Related Work

To our knowledge, four end-to-end pipelines have been proposed for LLP, with deep neural networks as the backbone architecture. Specifically, DLLP [Ardehaly and Culotta, 2017] is the first end-to-end LLP algorithm, with the KL divergence of prior and posterior proportions as the objective. Although DLLP can learn a competent instance-level classifier, it is hardly in accordance with the proportions in training data, especially with large bag sizes (e.g.,  $> 64$ ). In order to “guess” instance labels that are consistent to proportions [Yu *et al.*,

2013], [Dulac-Arnold *et al.*, 2019] perform convex relaxation to labeling variables and *first* recognize it as an OT problem. Aiming at end-to-end training, they instead impose inexact labeling with KL divergence similar to DLLP, and efficiently solve the resulting unbalanced OT problem with entropic regularization. However, the KL divergence results in inaccurate labeling, thus introduces additional noise to pseudo-labels compared to accurate labeling. We discuss the differences between their work and our work in Section 6.

Recently, LLP-GAN [Liu *et al.*, 2019] greatly improves DLLP with adversarial mechanism. In detail, the discriminator is designed as a  $(K+1)$ -way classifier, where the first  $K$  classes indicate the real ones, and the  $(K+1)^{th}$  class accounts for the generated one. The main insight is to produce better representation with adversarial learning, thus boosting the downstream discriminative tasks. In spite of substantial performance improvement compared with previous methods, LLP-GAN suffers from training instability, which inherits the characteristics from GANs [Gui *et al.*, 2020]. Besides, in order to obtain satisfactory results, it requires subtle network design and hyper-parameters selection. Similarly, based on visual adversarial training [Miyato *et al.*, 2018] in semi-supervised learning, [Tsai and Lin, 2020] introduce consistency regularization, and propose LLP-VAT, while the KL divergence on the proportions is intact in the loss.

However, KL divergence based methods hardly delivery classifiers abiding the proportions on training data. When performing alternate update, this inexact labeling exaggerates the label noise, especially with entropic regularization. In contrast, we address this problem with a constrained optimization to exactly follow the proportions. Furthermore, the diversity in unlabeled data demonstrates useful behavior for classification [Lu *et al.*, 2018]. Based on the great capacity of deep models, unsupervised representation learning is achievable and appealing as a promising step towards discriminative tasks. For example, aligning the clustering with semantic classes can successfully improve image classification and segmentation [Ji *et al.*, 2019]. [Asano *et al.*, 2020] apply (entropic constrained) optimal transport (OT) to learn representation with a self-labeling mechanism. Following their seminal framework, in this paper, we propose an OT based LLP algorithm, and further develop detailed techniques.

## 3 Preliminaries

In order to clearly describe how to leverage OT for LLP, we first introduce several preliminaries for both OT and LLP. More details are in Appendix due to the space limitation.

### 3.1 Optimal Transport via Kantorovich Relaxation

Define two discrete measures on two arbitrary sets  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively:  $\alpha = \sum_{i=1}^n a_i \delta_{x_i}$ ,  $\beta = \sum_{j=1}^m b_j \delta_{y_j}$ , where  $\mathbf{a} = (a_i)_{i=1}^n \in \Sigma^n$ ,  $\mathbf{b} = (b_j)_{j=1}^m \in \Sigma^m$ , and  $\Sigma^s$  is probability simplex, i.e.,  $\Sigma^s = \{\mathbf{a} \in \mathbb{R}_+^s \mid \sum_{i=1}^s a_i = 1\}$ ,  $\forall s \in \mathbb{Z}_+$ . To achieve mass splitting in *Kantorovich relaxation*, instead of permutation  $\sigma$  or surjective map  $T$ , we use *coupling matrix*  $\mathbf{P} \in \mathbb{R}_+^{n \times m}$  with  $P_{ij}$  as the mass flowing from  $x_i$  to  $y_j$ . Admissible couplings give a simpler characterization than maps

in Monge problem: Suppose  $\mathbf{a} \in \Sigma^n, \mathbf{b} \in \Sigma^m$ , we denote  $U(\mathbf{a}, \mathbf{b}) = \{\mathbf{P} \in \mathbb{R}_+^{n \times m} \mid \mathbf{P}\mathbf{1}_m = \mathbf{a}, \mathbf{P}^\top \mathbf{1}_n = \mathbf{b}\}$ , where  $\mathbf{1}_n$  represents  $n$ -dimensional all ones column vector. Hence,  $U(\mathbf{a}, \mathbf{b})$  stands for the set of all Kantorovich mappings. Similar to Monge problem, let  $\langle \cdot, \cdot \rangle$  be the Frobenius dot-product, given  $\mathbf{a}, \mathbf{b} \in \Sigma^n$  and a cost matrix  $\mathbf{C}$ , the Kantorovich's *optimal transport* (OT) problem is a linear programming as

$$\mathbf{L}_{\mathbf{C}}(\mathbf{a}, \mathbf{b}) = \min_{\mathbf{P} \in U(\mathbf{a}, \mathbf{b})} \langle \mathbf{P}, \mathbf{C} \rangle = \sum_{i,j} P_{ij} C_{ij}. \quad (1)$$

### 3.2 Entropic Regularization of Optimal Transport

The solution of the original OT problem (1) is non-unique and tends to be sparse, i.e., arriving at certain vertex of the polytope  $U(\mathbf{a}, \mathbf{b})$ . To form a more ‘‘blurred’’ prediction, [Cuturi, 2013] propose the entropic regularization for OT problem. To be specific, the discrete entropy of a coupling matrix  $\mathbf{P}$  is well-known as  $\mathbf{H}(\mathbf{P}) = -\sum_{i,j} P_{ij} \log(P_{ij})$ .

**Remark 1** ([Peyré et al., 2019]). *The entropic function  $\mathbf{H}(\cdot)$  is 1-strongly concave, due to the negative definite Hessian matrix:  $\partial^2 \mathbf{H}(\mathbf{P}) = -\text{diag}(1./\mathbf{P})$  and  $P_{ij} \leq 1, \forall i, j$ .*

Note that the rank-1 matrix  $\mathbf{a}\mathbf{b}^\top$  is an admissible coupling, and  $\mathbf{H}(\mathbf{P}) \leq \mathbf{H}(\mathbf{a}) + \mathbf{H}(\mathbf{b}) = \mathbf{H}(\mathbf{a}\mathbf{b}^\top)$ . The *entropic regularization* of OT adds  $-\mathbf{H}(\cdot)$  to original OT (1) as:

$$\mathbf{L}_{\mathbf{C}}^\varepsilon(\mathbf{a}, \mathbf{b}) = \min_{\mathbf{P} \in U(\mathbf{a}, \mathbf{b})} \langle \mathbf{P}, \mathbf{C} \rangle - \varepsilon \mathbf{H}(\mathbf{P}), \quad (2)$$

which is a constrained minimization problem of an  $\varepsilon$ -convex function (Remark 1), thus has a unique minimum. More importantly, the entropic constraint guarantees a computationally efficient process to find the solution, as a consequence of restricting the search for low cost joint probabilities within sufficient smooth tables [Cuturi, 2013]. Apparently, if  $\varepsilon \rightarrow +\infty$ , the solution of entropic regularized OT (2) is  $\mathbf{a}\mathbf{b}^\top$ . In addition, the solution of (2) converges to that of the original OT (1) as  $\varepsilon \rightarrow 0$ .

### 3.3 Learning from Label Proportions

In LLP problem, because the label proportions are available, we can restrict the instance-level self-labeling procedure with these proportional information using an OT framework. Before further discussion, we first give the formal formulation for LLP by directly considering a multi-class problem with  $K$  classes in this paper. With no prior knowledge, we further suppose that the training data consist of  $N$  *randomly generated* disjoint bags. Consequently, the training data can be expressed as  $\mathcal{D} = \{(\mathcal{B}_i, \mathbf{p}_i)\}_{i=1}^m$ , where  $\mathcal{B}_i = \{\mathbf{x}_{i,j}\}_{j=1}^{n_i}$  denotes the instances in the  $i^{\text{th}}$  bag, and  $\mathcal{B}_i \cap \mathcal{B}_j = \emptyset, \forall i \neq j$ . The  $\mathbf{p}_i \in [0, 1]^K$  and  $n_i$  are the known ground-truth label proportions and the bag size of the  $i^{\text{th}}$  bag, respectively.

## 4 Approach

### 4.1 Linking OT to LLP (PLOT)

[Asano et al., 2020] introduce a self-labeling framework to leverage equal clustering to learn discriminative representation on unsupervised data and achieve classification. In LLP problem, although the class distribution is not uniform within each bag, we can easily modify the admissible couplings in

(1) to fit in the proportional information. Consequently, with  $p_i^y$  as the proportion of class  $y$  in bag  $i$ , we have

$$\begin{aligned} \min_q BCE(p, q) &= -\sum_{i=1}^m \sum_{j=1}^{n_i} \sum_{y=1}^K \frac{q(y|\mathbf{x}_{i,j})}{n_i} \log p_\phi(y|\mathbf{x}_{i,j}) \\ \text{s.t.} \sum_{j=1}^{n_i} q(y|\mathbf{x}_{i,j}) &= p_i^y \cdot n_i, q(y|\cdot) \in [0, 1], \\ \forall y \in [K] &= \{1, 2, \dots, K\}, \forall i \in [m] = \{1, 2, \dots, m\}. \end{aligned} \quad (3)$$

Note that (3) is a constrained optimization and the labels should strictly comply with the proportions in each bag. Nevertheless, (3) will be combinatorial if imposing one-hot labeling in  $q$ , thus is very difficult to optimize. Fortunately, as we will point out, (3) is a typical OT problem, thus can be solved relatively efficiently to arrive at a sparse (one-hot) solution.

For better explanation, we rewrite (3) in a matrix fashion. Formally, let  $\mathbf{Q}^i = (Q_{jk}^i) \in \mathbb{R}_+^{K \times n_i}, Q_{jk}^i = q(k|\mathbf{x}_{i,j})/n_i$ , and  $\mathbf{P}^i = (P_{jk}^i) \in \mathbb{R}_+^{K \times n_i}, P_{jk}^i = p_\phi(k|\mathbf{x}_{i,j})/n_i$ . In addition, denote  $\mathbf{Q} = \text{diag}\{\mathbf{Q}^i\}_{i=1}^m, \mathbf{P} = \text{diag}\{\mathbf{P}^i\}_{i=1}^m$ , and  $\mathbf{p} = (\mathbf{p}_1^\top, \mathbf{p}_2^\top, \dots, \mathbf{p}_m^\top)^\top, \mathbf{b} = (\mathbf{1}_{n_1}^\top/n_1, \mathbf{1}_{n_2}^\top/n_2, \dots, \mathbf{1}_{n_m}^\top/n_m)^\top$ . Define  $U(\mathbf{p}, \mathbf{b}) = \{\mathbf{Q} \in \mathbb{R}_+^{m \times K \times N} \mid \mathbf{Q}\mathbf{1}_N = \mathbf{p}, \mathbf{Q}^\top \mathbf{1}_{mK} = \mathbf{b}\}$ . Accordingly, we have an equivalent OT problem for (3) as:

$$\min_{\mathbf{Q} \in U(\mathbf{p}, \mathbf{b})} \langle \mathbf{Q}, -\log \mathbf{P} \rangle = BCE(p, q) + \log \prod_{i=1}^m n_i. \quad (4)$$

On the other hand, with  $\lambda \rightarrow +\infty$ , we can instead solve the entropic regularized OT problem to accelerate the process of convergence, as well as attaining unique non-sparse solution.

$$\mathbf{L}_{-\log \mathbf{P}}^{1/\lambda}(\mathbf{p}, \mathbf{b}) = \min_{\mathbf{Q} \in U(\mathbf{p}, \mathbf{b})} \langle \mathbf{Q}, -\log \mathbf{P} \rangle - \frac{1}{\lambda} \mathbf{H}(\mathbf{Q}). \quad (5)$$

### 4.2 Alternating Optimization

In the proposed learning framework, the network parameters  $\phi = (\varphi, \theta)$  and self-labels  $\mathbf{Q}$  are alternately updated. Now, we further describe the details as follows.

**Training the Network with Fixed  $\mathbf{Q}$ .** Because the cross-entropy is differentiable w.r.t. the network parameters  $\phi = (\varphi, \theta)$ , we can directly conduct common optimizer, e.g., Adam, on the objective in (3) by fixing  $\mathbf{Q}$ .

**Updating the Labels  $\mathbf{Q}$  with Fixed  $\phi$ .** When the model is fixed, the label assignment matrix  $\mathbf{Q}$  are obtained by OT or entropic regularized OT. When performing original OT, the solution  $\mathbf{Q}^*$  is with 0/1 binary elements. However, entropic regularized OT produces  $\mathbf{Q}^*$  with all elements in  $[0, 1]$ . In practice, we employ two strategies for label update: hard labeling and soft labeling. In hard labeling, we update  $\mathbf{Q}$  as:

$$Q_{js}^i = \begin{cases} 1, & \text{if } s = \underset{k}{\text{argmax}} q(k|\mathbf{x}_{i,j}) \\ 0, & \text{otherwise} \end{cases}, i = 1, 2, \dots, m. \quad (6)$$

In soft labeling, we directly use the labels obtained by entropic regularized OT. In the experimental part, we provide the performance comparison on hard and soft labeling.

### 4.3 Two-stage Training for LLP

In practice, we consider to perform the clustering in every single bag, with the proportions as the constraint for instances number in each cluster. In detail, we conduct the constrained OT problem (4) w.r.t.  $\mathbf{Q}^i$  and  $\mathbf{P}^i$ , with minor revision on  $U(\mathbf{p}, \mathbf{b})$ , i.e.,  $U(\mathbf{p}_i, \mathbf{b}_i)$ , where  $\mathbf{b}_i = \mathbf{1}_{n_i}/n_i$ . On the other hand, we can instead solve the entropic regularized OT problem (5) with the same revision as (7) to accelerate the training, as well as obtain non-sparse solution to perform soft labeling:

$$\mathbf{L}_{-\log \mathbf{P}}^{1/\lambda}(\mathbf{p}_i, \mathbf{b}_i) = \min_{\mathbf{Q}^i \in U(\mathbf{p}_i, \mathbf{b}_i)} \langle \mathbf{Q}^i, -\log \mathbf{P} \rangle - \frac{1}{\lambda} \mathbf{H}(\mathbf{Q}^i). \quad (7)$$

Our two-stage training process uses KL-divergence based LLP training as the first stage, and supervised learning with pseudo-labels generated by OT as the second stage. In order to reduce the memorization of corrupt pseudo-labels, we incorporate symmetric cross-entropy (SCE) [Wang *et al.*, 2019] in second stage. Let  $L_{CE}$  be cross-entropy  $H(\mathbf{p}, \mathbf{q})$ , we denote  $L_{RCE}$  as reverse cross-entropy  $H(\mathbf{q}, \mathbf{p})$  defined as:

$$L_{RCE} = H(\mathbf{q}, \mathbf{p}) = - \sum_{k=1}^K p(k|x) \log q(k|x). \quad (8)$$

Accordingly, the SCE is defined as

$$SCE = L_{CE} + L_{RCE} = H(\mathbf{p}, \mathbf{q}) + H(\mathbf{q}, \mathbf{p}). \quad (9)$$

To further improve the performance, we apply mixup [Zhang *et al.*, 2018] to any two pairs  $(\mathbf{x}_i, \mathbf{y}_i)$  and  $(\mathbf{x}_j, \mathbf{y}_j)$ , i.e.,  $\mathbf{x} = \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_j$  and  $\mathbf{y} = \lambda \mathbf{y}_i + (1 - \lambda) \mathbf{y}_j$ , to obtain a new data point  $(\mathbf{x}, \mathbf{y})$ . The detailed training process of the second stage is shown in Algorithm 1.

## 5 Numerical Experiments

In order to demonstrate OL-LLP is model-agnostic, in this section, we conduct quantitative and qualitative studies with extensive experiments, to show the improvement of former LLP methods, when using PLOT as the second training stage. We use two benchmark datasets: CIFAR-10 and CIFAR-100. The comparisons are performed on three recently proposed algorithms DLLP [Ardehaly and Culotta, 2017], LLP-VAT [Tsai and Lin, 2020], and LLP-GAN [Liu *et al.*, 2019].

### 5.1 Synthetic Datasets Study

We first evaluate our approach with a toy example. In detail, we choose the well-known ‘‘two-moon’’ dataset to obtain the performance of DLLP and LLP-GAN, as well as their combination with PLOT as the second training stage. A 3-hidden-layer neural network with ReLU activation is employed. We conduct the experiment with the same bag setting, where 40 bags are generated with each containing 50 points.

The results are visualized in Figure 2, with different classes in red and green, respectively. For DLLP in Figure 2(a), blobs of data is mis-classified for both categories, while the portion of errors reduces when incorporating PLOT to DLLP in Figure 2(b). When it comes to stronger baseline LLP-GAN in Figure 2(c), only the tails of both categories are mis-classified. However, only visible two green points are mis-classified to red at the tail in Figure 2(d), whose results are obtained by conducting OP-LLP training as the second stage to

### Algorithm 1 Second Stage Pseudo Labeling with OT (PLOT)

#### Require:

The LLP training data  $\mathcal{D} = \{(\mathbf{B}_i, \mathbf{p}_i)\}_{i=1}^m$ ;  
 The initialization of the network parameters:  $\phi = (\varphi, \theta)$ ;  
 The block diagonal matrix  $\mathbf{Q} = \text{diag}\{\mathbf{Q}^i\}_{i=1}^m$  as the pseudo-labels obtained in the first-stage.

#### Ensure:

- 1: **for** each epoch over  $\mathcal{D}$  and the pseudo-labels  $\mathbf{Q}$  **do**
- 2:   Sample one mini-batch of training bags  $B_D$  from  $\mathcal{D}$  and the corresponding pseudo-labels from  $\mathbf{Q}$
- 3:   **for** each mini-batch **do**
- 4:     Train network based on SCE with the pseudo-labels to update the parameters  $\phi = (\varphi, \theta)$ .
- 5:   **end for**
- 6:   Update the pseudo labeling matrix  $\mathbf{Q} = \text{diag}\{\mathbf{Q}^i\}_{i=1}^m$  through the OT or entropic regularized OT.
- 7: **end for**
- 8: **return** The appropriate values of  $\phi = (\varphi, \theta)$ .

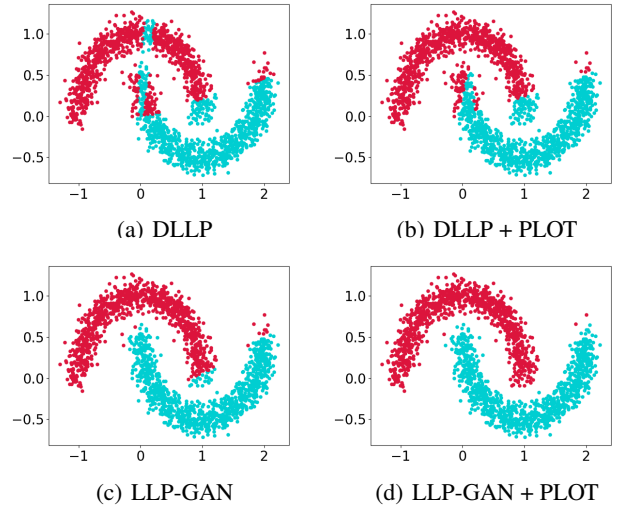


Figure 2: Comparison of DLLP, LLP-GAN and their combinations with PLOT on the ‘‘two moons’’ dataset. (Best viewed in color)

LLP-GAN. It effectively demonstrates that our second stage can further boost the performance of other LLP solvers. In practice, solving OT can be regarded as a complementary refinement for the solution of minimizing the KL divergence.

### 5.2 Experimental Setting

**Label Proportions Generation.** As there is no off-the-shelf LLP datasets, we first generate LLP datasets with the label proportions using four supervised benchmark datasets. Following the setting from [Liu *et al.*, 2019], we construct four kinds of bags, with bag sizes of 16, 32, 64, and 128, respectively. In order to avoid the influence of different label distributions, the bag setting is fixed across different methods.

**Training Setting.** We choose a conv-based 13-layer max-pooling network as the backbone for all methods. The architecture details is given in Appendix A.6. Meanwhile, Adam

optimizer is used with  $\beta_1 = 0.5$  and  $\beta_2 = 0.999$ . The initial learning rate is  $1e-4$ , divided by 2 every 100 epochs. Random horizontal flip and random crop with padding to the original images is employed as data augmentation. The hyperparameters for SCE is the same as [Wang *et al.*, 2019].

### 5.3 Comparison with the State-of-the-arts

In this section, we provide the overall accuracy comparison between our approach and three SoTA methods: DLLP, LLP-GAN, and LLP-VAT, on CIFAR-10 and CIFAR-100. Specifically, our method divides the training into two phases, where the first phase is to train the three KL-divergence based models as the *teacher model*, and the second stage is to update the *student model* based on the pseudo-labels obtained in the first step, followed by solving the OT with the proportional information in an alternate manner, described in Algorithm 1. In practice, note that our approach is model-agnostic, which means the first stage can be any deep learning based model, and our model can boost their performance by attaining local solutions with significantly better performance.

For fair comparison, in Tabel 1 we provide the results on CIFAR-10 without mixup [Zhang *et al.*, 2018]. For the first glance, since bigger bag size contains less supervision when the training data is fixed, the accuracies of all the methods decrease along with the increase of bag size. Nevertheless, the performance of all methods can efficiently improve when combining with the proposed second-stage training. Note that the improvement is more obvious with bag sizes of 64 and 128, which is more challenging compared with smaller bag sizes. For example, the accuracy of LLP-GAN + PLOT with bag size 128 reaches 79.08, even superior to the previous SoTA result with bag size 64. Furthermore, this combination can approximately achieve a similar performance to the fully supervised scenario with small bag sizes, e.g., 16. In general, we can further draw a conclusion: Algorithm with the best performance in the first stage consistently results in the best performance when introducing PLOT as the second stage.

When it comes to harder CIFAR-100 in Table 2, we have similar results: Our method can considerably boost other LLP methods to reach new state-of-the-art results. In particular, our method reaches an accuracy rate of 43.44 with bag size 128, which is much higher than all the previous state-of-the-arts. This indicates that incorporating the OT based pseudo-labels mechanism to learn label proportions is of higher efficiency. In practice, the proportion of pseudo labeling can be intact to the real proportions through the optimal transport, thus improving the total accuracy of training data.

Note that the improvement for LLP-VAT is relatively limited compared to DLLP and LLP-GAN when incorporating our post-hoc training, especially for bag sizes of 16 and 32. A main reason is the key idea for noise reduction in network prediction is coincident in both methods, although different mechanisms are applied: LLP-VAT leverages the prediction consistency with adversarial samples as regularization, while our recipe is to obtain pseudo-labels with OT, then employs SCE and mixup as the post-hoc noise reduction techniques.

The motivation behind our approach is the predicted labels of training data can be further refined by optimal transport, so as to fulfill better pseudo labeling. To show this, we provide




Figure 3: Comparison of accuracy on CIFAR-10 and CIFAR-100 in the training process. (Best viewed in color)

the accuracy curves for CIFAR-10 and CIFAR-100 with bag size 64 in Figure 3, where TD, TD\_OT, and TD\_OT\_5 denote the accuracies without OT, with OT, and with the average of the last five pseudo-labels with OT (i.e., ensembles), respectively. It demonstrates that the accuracy can significantly improve after involving optimal transport. In particular, the advantage is more obvious for CIFAR-100, by a larger margin between TD and TD\_OT. Meanwhile, the accuracy can further improve by using the average of the last five OT pseudo-labels as the final pseudo-label in every OT labeling phase.

### 5.4 Ablation Study

In this section, we carry out ablation study on key components introduced in PLOT. All the experiments are conducted on the dataset CIFAR-10 with the bag size 64. More specifically, our analysis mainly focuses on the difference between cross-entropy and symmetric cross-entropy. Furthermore, we evaluate the effect of mixup augmentation to our model when updating network parameters with above losses.

The final results are shown in Table 3. The baseline is the performance of DLLP in the first stage. The  $L_{CE}$  and  $L_{SCE}$  means that the model is trained without optimal transport, merely using the mean of instance-level predictions in the previous five epochs as pseudo-labels. This strategy can also improve the performance slightly. In particular, compared with  $L_{CE}$ ,  $L_{SCE}$  provides a 2% higher performance improvement. On the other hand, the accuracy can be further increased when employing the pseudo-labels generated by OT. In particular, SCE is considerably effective, in contrast to CE, mainly due to its robustness to label noise.

Besides, mixup can further improve the performance significantly, which is mainly because of the following two reasons. On one hand, it reduces the memorization of corrupt labels, similar to SCE. On the other hand, it can impose more generalization on the network with the convex combination of a pair of examples, as well as their pseudo-labels.

### 5.5 Hard-label vs. Soft-label

In our model, we can employ two strategies to update the label: the hard labeling with the sparse original OT solution and the soft labeling with entropic regularized OT solution. The detailed process is shown in (6). In order to better demonstrate the difference between the hard-labels and soft-labels in our approach, we adopt CIFAR-10 and CIFAR-100 to con-

Model	Bag Size			
	16	32	64	128
DLLP [Ardehaly and Culotta, 2017]	87.69 (0.54)	82.88 (0.47)	71.11 (0.56)	46.68 (0.75)
DLLP + PLOT ( <b>ours</b> )	<b>90.01</b> (0.23)	86.87 (0.24)	79.69 (0.44)	55.21 (0.56)
LLP-VAT [Tsai and Lin, 2020]	88.36 (0.29)	83.01 (0.46)	70.53 (0.98)	51.25 (0.88)
LLP-VAT + PLOT ( <b>ours</b> )	89.87 (0.31)	87.89 (0.29)	79.33 (0.39)	63.28 (0.43)
LLP-GAN [Liu <i>et al.</i> , 2019]	86.97 (0.42)	83.13 (0.43)	77.21 (0.39)	68.37 (1.21)
LLP-GAN + PLOT ( <b>ours</b> )	89.15 (0.35)	<b>88.21</b> (0.23)	<b>84.14</b> (0.68)	<b>79.09</b> (0.44)

Table 1: Test accuracy rates and standard deviations (%) on CIFAR10 with different bag sizes. The results are obtained for 5 runs.

Model	Bag Size			
	16	32	64	128
DLLP [Ardehaly and Culotta, 2017]	61.95 (0.54)	55.22 (0.77)	41.21 (1.36)	15.17 (1.07)
DLLP + PLOT ( <b>ours</b> )	65.20 (0.41)	61.37 (0.74)	49.55 (0.86)	23.10 (0.29)
LLP-VAT [Tsai and Lin, 2020]	65.21 (0.39)	58.65 (0.43)	45.53 (1.09)	21.25 (0.81)
LLP-VAT + PLOT ( <b>ours</b> )	65.39 (0.28)	<b>62.01</b> (0.30)	52.75 (0.51)	28.23 (0.4)
LLP-GAN [Liu <i>et al.</i> , 2019]	61.66 (0.49)	56.78 (0.55)	50.29 (1.12)	33.65 (1.25)
LLP-GAN + PLOT ( <b>ours</b> )	<b>65.41</b> (0.35)	61.68 (0.48)	<b>55.66</b> (0.41)	<b>43.44</b> (0.71)

Table 2: Test accuracy rates and standard deviations (%) on CIFAR100 with different bag sizes. The results are obtained for 5 runs.

Method	CIFAR-10
baseline	71.11 (0.36)
$L_{CE}$	72.27 (0.39)
$L_{SCE}$	74.21 (0.24)
$L_{CE} + OT$	78.57 (0.32)
$L_{SCE} + OT$	79.69 (0.44)
$L_{CE} + OT + MIXUP$	84.39 (0.22)

 Table 3: Ablation study on CIFAR-10 with bag size 64:  $L_{CE}$  (cross-entropy),  $L_{SCE}$  (symmetric cross-entropy); OT (update pseudo-labels with optimal transport), MIXUP (mixup data augmentation).

duct comparison on the performance with different bag sizes. Specifically, the solution of the first stage is fixed for both hard and soft labeling for fair comparison. The results are shown in Figure 4, where we provide the convergence curve of the second stage under different pseudo labelings.

From the result, we observe that it reach a comparable performance with hard and soft labelings for CIFAR-10 and CIFAR-100. Intuitively, the soft-labels are more informative than the hard labels. However, it does not necessarily lead to better performance due to the potential noisy labeling. On the other hand, the solution of the original OT tends to be sparse, and with less noise. We reckon this as the main reason of the comparable performance of two psuedo labeling strategies.

## 6 Discussion and Conclusion

In this paper, we study the challenge in existing LLP methods, and point out the minimization on KL divergence between the prior and posterior class distributions is inadequate to comply with the label proportions exactly. From this perspective, we propose a second stage for existing LLP solver with a framework to combine instance-level classification and pseudo labeling, and alternately optimize these two objectives.

Compared with existing LLP solvers, especially [Dulac-




Figure 4: Comparison of hard and soft pseudo labelings on two benchmarks. (Best viewed in color)

Arnold *et al.*, 2019], our improvements are four-fold. First, we introduce exact pseudo labeling to convert the unconstrained KL divergence minimization into a constrained one, so that the labeling can strictly meet the proportions and avoid suboptimal solutions. Second, although the pseudo-labels can be efficiently found by Sinkhorn’s algorithm for entropic regularized OT (or linear programming for OT), we instead recognize the resulting supervised learning with pseudo-labels as a noisy label problem, due to the uncertainty introduced by the proportional information. Accordingly, we apply mix-up data augmentation and use SCE instead of CE as the differentiable loss function, which can be minimized with any line search optimizer (e.g., Adam). Third, we propose a two-stage LLP training by integrating PLOT as the second phase, and elaborately demonstrate that our framework can further improve the performance of DLLP, LLP-VAT, and LLP-GAN with extensive experiments, thus is model-agnostic. Lastly, instead of directly using the pseudo-labels, we empirically study the difference between hard and soft labeling strategies, and provide suggestions for practical usage.

## References

- [Ardehaly and Culotta, 2017] Ehsan Mohammady Ardehaly and Aron Culotta. Co-training for demographic classification using deep learning from label proportions. In *the IEEE International Conference on Data Mining Workshops*, pages 1017–1024, 2017.
- [Asano *et al.*, 2020] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations*, 2020.
- [Caron *et al.*, 2018] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *the European Conference on Computer Vision*, pages 132–149, 2018.
- [Chen *et al.*, 2014] Tao Chen, Felix X Yu, Jiawei Chen, Yin Cui, Yan-Ying Chen, and Shih-Fu Chang. Object-based visual sentiment concept analysis and application. In *the 22nd ACM International Conference on Multimedia*, pages 367–376, 2014.
- [Cuturi, 2013] Marco Cuturi. Sinkhorn distances: Light-speed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.
- [Dulac-Arnold *et al.*, 2019] Gabriel Dulac-Arnold, Neil Zeghidour, Marco Cuturi, Lucas Beyer, and Jean-Philippe Vert. Deep multi-class learning from label proportions. *arXiv preprint arXiv:1905.12909*, 2019.
- [Gui *et al.*, 2020] Jie Gui, Zhenan Sun, Yonggang Wen, Dacheng Tao, and Jieping Ye. A review on generative adversarial networks: Algorithms, theory, and applications. *arXiv preprint arXiv:2001.06937*, 2020.
- [Hernández-González *et al.*, 2018] Jerónimo Hernández-González, Inaki Inza, Lorena Crisol-Ortíz, María A Guembe, María J Iñarra, and Jose A Lozano. Fitting the data from embryo implantation prediction: Learning from label proportions. *Statistical Methods in Medical Research*, pages 1056–1066, 2018.
- [Ji *et al.*, 2019] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *the IEEE International Conference on Computer Vision*, pages 9865–9874, 2019.
- [Kuck and de Freitas, 2012] Hendrik Kuck and Nando de Freitas. Learning about individuals from group statistics. *arXiv preprint arXiv:1207.1393*, 2012.
- [Lai *et al.*, 2014] Kuan-Ting Lai, Felix X Yu, Ming-Syan Chen, and Shih-Fu Chang. Video event detection by inferring temporal instance labels. In *the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2243–2250, 2014.
- [Liebig *et al.*, 2015] Thomas Liebig, Marco Stolpe, and Katharina Morik. Distributed traffic flow prediction with label proportions: from in-network towards high performance computation with MPI. In *the International Conference on Mining Urban Data-Volume 1392*, pages 36–43. CEUR-WS. org, 2015.
- [Liu *et al.*, 2019] Jiabin Liu, Bo Wang, Zhiqian Qi, Yingjie Tian, and Yong Shi. Learning from label proportions with generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 7167–7177, 2019.
- [Lu *et al.*, 2018] Nan Lu, Gang Niu, Aditya Krishna Menon, and Masashi Sugiyama. On the minimal supervision for training any binary classifier from only unlabeled data. In *International Conference on Learning Representations*, 2018.
- [Miyato *et al.*, 2018] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):1979–1993, 2018.
- [Peyré *et al.*, 2019] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- [Sun *et al.*, 2017] Tao Sun, Dan Sheldon, and Brendan O’Connor. A probabilistic approach for learning with label proportions applied to the US presidential election. In *the IEEE International Conference on Data Mining*, pages 445–454. IEEE, 2017.
- [Tsai and Lin, 2020] Kuen-Han Tsai and Hsuan-Tien Lin. Learning from label proportions with consistency regularization. In *Asian Conference on Machine Learning*, pages 513–528. PMLR, 2020.
- [Villani, 2008] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [Wang *et al.*, 2019] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *the IEEE International Conference on Computer Vision*, pages 322–330, 2019.
- [Yu *et al.*, 2013] Felix X Yu, Dong Liu, Sanjiv Kumar, Tony Jebara, and Shih-Fu Chang.  $\infty$ SVM for learning with label proportions. In *the International Conference on Machine Learning*, 2013.
- [Yu *et al.*, 2014a] Felix X Yu, Liangliang Cao, Michele Merler, Noel Codella, Tao Chen, John R Smith, and Shih-Fu Chang. Modeling attributes from category-attribute proportions. In *the 22nd ACM International Conference on Multimedia*, pages 977–980, 2014.
- [Yu *et al.*, 2014b] Felix X Yu, Krzysztof Choromanski, Sanjiv Kumar, Tony Jebara, and Shih-Fu Chang. On learning from label proportions. *arXiv preprint arXiv:1402.5902*, 2014.
- [Zhang *et al.*, 2018] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018.