Deep Partial Multi-Label Learning with Graph Disambiguation

Haobo Wang[1] , **Shisong Yang**[2] , **Gengyu Lyu**[2*] , **Weiwei Liu**[3] , **Tianlei Hu**[1] ,
Ke Chen[1] , **Songhe Feng**[4] and **Gang Chen**[1]

[1]Zhejiang University
[2]Beijing University of Technology
[3]Wuhan University
[4]Beijing Jiaotong University

{wanghaobo, htl, chenk, cg}@zju.edu.cn, yangshisong@emails.bjut.edu.cn,
lyugengyu@bjut.edu.cn, liuweiwei863@gmail.com, shfeng@bjtu.edu.cn

Abstract

In partial multi-label learning (PML), each data example is equipped with a candidate label set, which consists of multiple ground-truth labels and other false-positive labels. Recently, graph-based methods, which demonstrate a good ability to estimate accurate confidence scores from candidate labels, have been prevalent to deal with PML problems. However, we observe that existing graph-based PML methods typically adopt linear multi-label classifiers and thus fail to achieve superior performance. In this work, we attempt to remove several obstacles for extending them to deep models and propose a novel deep **P**artial multi-**L**abel model with gr**A**ph-disamb**I**guatio**N** (**PLAIN**). Specifically, we introduce the instance-level and label-level similarities to recover label confidences as well as exploit label dependencies. At each training epoch, labels are propagated on the instance and label graphs to produce relatively accurate pseudo-labels; then, we train the deep model to fit the numerical labels. Moreover, we provide a careful analysis of the risk functions to guarantee the robustness of the proposed model. Extensive experiments on various synthetic datasets and three real-world PML datasets demonstrate that PLAIN achieves significantly superior results to state-of-the-art methods.

1 Introduction

The rapid growth of deep learning allows us to tackle increasingly sophisticated problems. Among them, multi-label learning (MLL), which assigns multiple labels to an object, is a fundamental and intriguing task in various real-world applications such as image retrieval [Lai *et al.*, 2016], autonomous vehicles [Chen *et al.*, 2019a], and sentiment analysis [Yilmaz *et al.*, 2021]. Typically, these tasks require precisely labeled data, which is expensive and time-consuming due to the complicated structure and the high volume of the output space.

To address this issue, plenty of works have studied different weakly-supervised settings of MLL, including semi-supervised MLL [Shi *et al.*, 2020], multi-label with missing labels [Durand *et al.*, 2019; Huynh and Elhamifar, 2020], and single-labeled MLL [Xu *et al.*, 2022]. While these learning paradigms mainly battle the rapid growth of data volume, they overlook the intrinsic difficulty and ambiguity in multi-label annotation. For example, it is hard for a human annotator to identify an Alaskan Malamute from a Siberian Husky. One potential solution is to randomly decide the binary label value. However, it may involve uncontrollable label noise that has a negative impact on the training procedure. Another strategy is to leave it missing, but then, we may lose critical information. Recently, partial multi-label learning (PML) [Xie and Huang, 2018] has been proposed to alleviate this problem. In PML, the annotators are allowed to provide a candidate label set that contains all potentially relevant labels as well as false-positive labels.

A straightforward approach to learning from PML data is to regard all candidate labels as valid ones, and then, off-the-shelf multi-label methods can be applied. Nevertheless, the presence of false-positive labels can highly impact predictive performance. To cope with this issue, many PML methods [Xie and Huang, 2018; Yu *et al.*, 2018; He *et al.*, 2019; Sun *et al.*, 2019; Li *et al.*, 2020; Lyu *et al.*, 2020; Yan *et al.*, 2021; Xie and Huang, 2022] have been developed. Amongst them, graph-based PML methods [Wang *et al.*, 2019; Li and Wang, 2020; Xu *et al.*, 2020; Lyu *et al.*, 2020; Zhang and Fang, 2021] have been popular. In practice, graph-based algorithms demonstrate a good ability to obtain relatively precise confidence of candidate labels by aggregating information from the nearest neighbors. However, we observe that most of them adopt linear classifier, which is powerless to deal with complicated tasks and high data volumes.

Therefore, there is an urgent need for developing an effective deep model to boost the performance of graph PML methods. However, there are two main challenging issues. First, existing graph-based PML methods typically require propagating labels over the whole training dataset. It restricts their scalability when employing deep classifiers since the deep models are trained in batch mode. Second, it is well acknowledged that there exist label correlations in multi-label

---

datasets, which also helps disambiguate the candidate labels. For example, if labels *Amusement Park* and *Mickey Mouse* are identified as true, it is very likely that label *Disney* is true as well. Thus, it is required to exploit the label correlations to improve the disambiguation ability in a unified framework.

In this work, we propose a novel deep **P**artial multi-**L**abel model with gr**A**ph disamb**I**guatio**N** (PLAIN). In PLAIN, we involve both instance- and label-level similarities to disambiguate the candidate labels set, which simultaneously leverages instance smoothness as well as label correlations. Then, we build a deep multi-label classifier to fit the relatively credible pseudo-labels generated through disambiguation. Instead of separately training in two stages like previous works [Wang *et al.*, 2019; Zhang and Fang, 2021], we propose an efficient optimization framework that iteratively propagates the labels and trains the deep model. Moreover, we give a careful analysis of the risk functions such that proper risk functions are chosen for improved robustness. Empirical results on nine synthetic datasets and three real-world PML datasets clearly verify the efficiency and effectiveness of our method. More theoretical and empirical results can be found in Appendix[1].

2   Related Works

Multi-Label Learning.   MLL [Liu *et al.*, 2022; Wei *et al.*, 2022] aims at assigning each data example multiple binary labels simultaneously. An intuitive solution, the one-versus-all method (OVA) [Zhang and Zhou, 2014], is to decompose the original problem into a series of single-labeled classification problems. However, it overlooks the rich semantic information and dependencies among labels, and thus fails to obtain satisfactory performance. To solve this problem, many MLL approaches are proposed, such as embedding-based methods [Yu *et al.*, 2014; Shen *et al.*, 2018], classifier chains [Read *et al.*, 2011], tree-based algorithms [Wei *et al.*, 2021], and deep MLL models [Yeh *et al.*, 2017; Chen *et al.*, 2019b; Bai *et al.*, 2020]. Nevertheless, these methods typically require a large amount of fully-supervised training data, which is expensive and time-consuming in MLL tasks. To this end, some works have studied different weakly-supervised MLL settings. For example, semi-supervised MLL [Niu *et al.*, 2019; Shi *et al.*, 2020] leverages information from both fully-labeled and unlabeled data. Multi-label with missing labels [Durand *et al.*, 2019; Wei *et al.*, 2018; Xu *et al.*, 2022] allows providing a subset of ground-truth. In noisy MLL [Veit *et al.*, 2017], the binary labels may be flipped to the wrong one. Nevertheless, though these settings relieve the burden of multi-label annotating, they ignore the inherent difficulty of labeling, i.e. the labels can be ambiguous.

Partial Multi-Label Learning.   PML learns from a superset of ground-truth labels, where multiple labels are possible to be relevant. To tackle the ambiguity in the labels, the pioneering work [Xie and Huang, 2018] estimates the confidence of each candidate label being correct via minimizing confidence-weighted ranking loss. Some works [Yu *et al.*, 2018; Sun *et al.*, 2019; Li *et al.*, 2020; Gong *et al.*, 2022] recover the ground-truth by assuming the true label matrix is

low-rank. Recently, graph-based methods [Wang *et al.*, 2019; Xu *et al.*, 2020; Lyu *et al.*, 2020; Xie and Huang, 2020; Zhang and Fang, 2021; Xu *et al.*, 2021; Tan *et al.*, 2022] have attracted much attention from the community due to the good disambiguation ability. For instance, PARTICLE [Zhang and Fang, 2021] identifies the credible labels through an iterative label propagation procedure and then, applies pair-wise ranking techniques to induce a multi-label predictor. Although graph-based models have shown promising results, we observe that most of them adopt linear classifiers, which significantly limits the model expressiveness. Besides, they typically separate the graph disambiguation and model training into two stages, which makes the multi-label classifier error-prone; in contrast, the proposed PLAIN is an end-to-end framework.

Some studies have also developed deep PML models, such as adversarial PML model [Yan and Guo, 2019] and mutual teaching networks [Yan *et al.*, 2021]. However, it remains urgent to combine the graph and deep models together such that the model enjoys high expressiveness as well as good disambiguation ability. It should also be noted that the recent popular graph neural networks (GNN) model [Wu *et al.*, 2021] is not suitable for PML. The reason is that GNN aims at better representation learning (at an instance-level), while in PML, the graph structure is used to alleviate the label ambiguity.

3   Method

We denote the d-dimensional feature space as $\mathcal{X} \subset \mathbb{R}_d$, and the label space as $\mathcal{Y} = \{1, 2, \ldots, L\}$ with L class labels. The training dataset $\mathcal{D} = \{(\boldsymbol{x}_i, S_i)|1 \leq i \leq n\}$ contains n examples, where $\boldsymbol{x}_i \in \mathcal{X}$ is the instance vector and $S_i \subset \mathcal{Y}$ is the candidate label set. Without loss of generality, we assume that the instance vectors are normalized such that $||\boldsymbol{x}_i||_2 = 1$. Besides, we denote $\tilde{S}_i \subset \mathcal{Y}$ as the ground-truth label set, which is a subset of S_i. For simplicity, we define $\boldsymbol{y}_i, \tilde{\boldsymbol{y}}_i \in \{0, 1\}^L$ as the binary vector representations of S_i and $\tilde{S}_i$. When considering the whole dataset, we also denote the candidate label matrix by $\boldsymbol{Y} = [\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n]^\top \in \mathbb{R}^{n \times L}$.

In what follows, we first introduce the learning target of our model that integrates label correlations to boost disambiguation ability. Then, we provide an efficient optimization framework for our deep propagation network.

3.1   Learning Objective

Our ultimate goal is to estimate a multi-label predictor $f : \mathcal{X} \mapsto [0, 1]^L$ from $\mathcal{D}$. In this work, we assume that its output $\hat{\boldsymbol{y}} = f(\boldsymbol{x})$ satisfies the following properties,

1. $\hat{\boldsymbol{y}}$ is close to the candidate vector $\boldsymbol{y}$;

2. $\hat{\boldsymbol{y}}$ satisfies the **instance-level similarity**. That is, if two instances are close to each other, then their label vectors are likely to be close too;

3. $\hat{\boldsymbol{y}}$ satisfies the **label-level similarity**. That is, if two labels co-occur in the training data frequently, then they tend to co-exist at many instances.

These properties coincide with our goals. The first property is natural since the only thing we can trust is the candidate

---
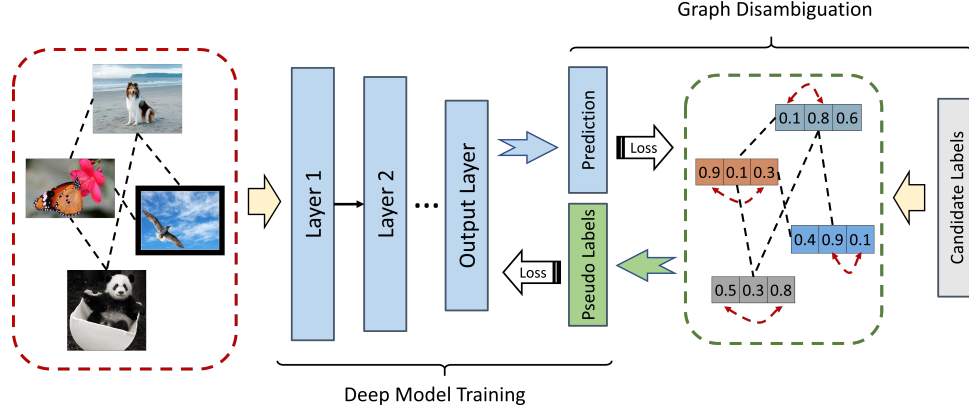[1]https://arxiv.org/abs/2305.05882

Figure 1: The model architecture of PLAIN. At each epoch, we first train the deep model by traveling the training dataset and then propagating the labels with the help of candidate labels and model prediction; PLAIN alternates between the two operations.

label set. The second one has been widely adopted to disambiguate the ground-truth label from the false-positive labels. The last one exploits the label correlations to further promote the performance. In this paper, we adopt a graph-based model to achieve these goals. Generally speaking, the graph structure can be collected in various ways. One direct strategy is to use prior knowledge. For instance, for data mining tasks, social networks [Qiu *et al.*, 2018] can be used to depict instance similarity. The label similarity can also be estimated using an external knowledge graph [Lee *et al.*, 2018].

However, we notice that there is no available graph structure in existing public PML datasets. Hence, we propose to construct the instance graph $\boldsymbol{A}^x \in \mathbb{R}^{n \times n}$ and the label graph $\boldsymbol{A}^y \in \mathbb{R}^{L \times L}$ in a data-driven way. To obtain $\boldsymbol{A}^x$, we adopt a monomial kernel [Iscen *et al.*, 2017] as the distance metric and calculate a sparse instance affinity matrix $\boldsymbol{S} \in \mathbb{R}^{n \times n}$ with elements,

$$s_{ij} := \begin{cases} \max(\boldsymbol{x}_i^\top \boldsymbol{x}_j, 0)^\rho, & \text{if } i \neq j \text{ and } j \in \mathcal{N}_k(\boldsymbol{x}_i) \\ 0, & \text{otherwise} \end{cases} .$$

Here $\mathcal{N}_k(\boldsymbol{x}_i)$ saves the indices of k-nearest neighbors of $\boldsymbol{x}_i$. $\rho > 0$ controls the weights of similarity. Then, we get a symmetrized instance graph by $\boldsymbol{A}^x = \boldsymbol{S} + \boldsymbol{S}^\top$.

Inspired by HALE [Lyu *et al.*, 2020], we build the label graph $\boldsymbol{A}^y$ by means of the label co-occurrence statistics from the training data, whose elements are given by,

$$a_{ij}^y := \frac{\sum_{k=1}^n \mathbb{I}(y_{ki} = 1 \wedge y_{kj} = 1)}{\sum_{k=1}^n \mathbb{I}(y_{ki} = 1) + \sum_{k=1}^n \mathbb{I}(y_{kj} = 1)},$$

where $\mathbb{I}(\cdot)$ is an indicator function returning 1 if the event is true and 0 otherwise. It is worth noting that HALE simply calculates the average co-occurrence by dividing the instance number, which overlooks the existence of false-positive labels. In our method, we penalize it by the counts of instances that are assigned the labels as the candidates. In other words, if one label is assigned to too many instances, its co-occurrence statistics are less reliable.

Next, we impose the instance-level similarity property on the model output by means of a Laplacian regularizer,

$$\mathcal{J}_x(\hat{\boldsymbol{Y}}) = \sum_{i,j=1}^n a_{ij}^x \left\| \frac{\hat{\boldsymbol{y}}_i}{\sqrt{d_i^x}} - \frac{\hat{\boldsymbol{y}}_j}{\sqrt{d_j^x}} \right\|^2 = \text{tr}(\hat{\boldsymbol{Y}}^\top \boldsymbol{L}_x \hat{\boldsymbol{Y}}), \quad (1)$$

where $\hat{\boldsymbol{Y}} = [\hat{\boldsymbol{y}}_1, \ldots, \hat{\boldsymbol{y}}_n]^\top$ is the predicted label matrix. $\boldsymbol{L}_x = \boldsymbol{I} - \boldsymbol{D}_x^{-\frac{1}{2}} \boldsymbol{A}^x \boldsymbol{D}_x^{-\frac{1}{2}}$ is the Laplacian matrix of instance graph. $\boldsymbol{D}_x = \text{diag}[d_1^x, \ldots, d_n^x]$ is a diagonal matrix with $d_i^x = \sum_{j=1}^n a_{ij}^x$. The objective in Eq. (1) enforces two data examples to have similar predicted labels if they have a large affinity score. Similarly, we can introduce the label-level similarity by,

$$\mathcal{J}_y(\hat{\boldsymbol{Y}}) = \text{tr}(\hat{\boldsymbol{Y}} \boldsymbol{L}_y \hat{\boldsymbol{Y}}^\top), \quad \boldsymbol{L}_y = \boldsymbol{I} - \boldsymbol{D}_y^{-\frac{1}{2}} \boldsymbol{A}^y \boldsymbol{D}_y^{-\frac{1}{2}}.$$

Finally, we involve a candidate consistency term to get the following objective,

$$\mathcal{J}(\hat{\boldsymbol{Y}}, \boldsymbol{Y}) = \frac{\eta}{2} \|\hat{\boldsymbol{Y}} - \boldsymbol{Y}\|_F^2 + \frac{\alpha}{2} \mathcal{J}_x(\hat{\boldsymbol{Y}}) + \frac{\beta}{2} \mathcal{J}_y(\hat{\boldsymbol{Y}}),$$

where $\|\cdot\|_F$ is the Frobenius norm. η, α, and β are all positive hyperparameters.

In summary, the main goal of our objective is to utilize the instance- and label-level similarity to recover reliable label confidences and address the problem of false positives, while simultaneously enjoying an extra performance improvement by exploiting label correlations.

3.2 Deep Propagation Network

If a linear model is adopted, the parameters can be directly obtained by optimizing $\mathcal{J}(f(\boldsymbol{X}), \boldsymbol{Y})$ with gradient descent. However, in deep models, it is not the case. When we update the model using a mini-batch of data, the instance-level manifold regularizer can involve unexpected gradient updating on the nearest neighbors of these data. Then, the training time will increase rapidly. To remedy this problem, we propose an efficient optimization framework for training a deep model from PML data. Specifically, we maintain a smooth pseudo-label matrix $\boldsymbol{Z}$, which minimizes the propagation error while being close to the model output $\hat{\boldsymbol{Y}}$. In the sequel, we elaborate the training procedure of our model. The training procedure is also summarized in Algorithm 1.

Propagation Step

At t-th epoch, we first propagate the label information from candidates $\boldsymbol{Y}$ and the model output $\hat{\boldsymbol{Y}}$ to the intermediate ones. In graph-based methods, the output label is required to be smooth enough. Therefore, we use the Frobenius norm as the measure to ensure the closeness of $\boldsymbol{Z}$ and $\hat{\boldsymbol{Y}}$. Formally, we estimate $\boldsymbol{Z}$ by,

$$\min_{\boldsymbol{Z}} \mathcal{L}_{\text{LP}}(\boldsymbol{Z}) = \frac{1}{2}||\boldsymbol{Z} - \hat{\boldsymbol{Y}}||_F^2 + \mathcal{J}(\boldsymbol{Z}, \boldsymbol{Y}). \quad (2)$$

We can derive its gradient by,

$$\frac{\partial \mathcal{L}_{\text{LP}}}{\partial \boldsymbol{Z}} = (1+\eta)\boldsymbol{Z} + \alpha \boldsymbol{L}_x \boldsymbol{Z} + \beta \boldsymbol{Z} \boldsymbol{L}_y - (\hat{\boldsymbol{Y}} + \eta \boldsymbol{Y}). \quad (3)$$

which is followed by the matrix properties that $\frac{\partial ||\boldsymbol{Z}-\boldsymbol{Y}||_F^2}{\partial \boldsymbol{Z}} = 2(\boldsymbol{Z} - \boldsymbol{Y})$ and $\frac{\partial \text{tr}(\boldsymbol{Z}^\top \boldsymbol{L}_x \boldsymbol{Z})}{\partial \boldsymbol{Z}} = (\boldsymbol{L}_x + \boldsymbol{L}_x^\top)\boldsymbol{Z}$. We have $\boldsymbol{L}_x + \boldsymbol{L}_x^\top = 2\boldsymbol{L}_x$ because $\boldsymbol{L}_x$ is symmetric. The optimal condition $\frac{\partial \mathcal{L}_{\text{LP}}}{\partial \boldsymbol{Z}} = \boldsymbol{0}$ typically requires solving a Sylvester system, which can be time-consuming. Therefore, we solve this problem by gradient descent $\boldsymbol{Z} = \boldsymbol{Z} - \gamma \frac{\partial \mathcal{L}_{\text{LP}}(\boldsymbol{Z})}{\partial \boldsymbol{Z}}$, where $\gamma > 0$ is the learning rate. Note that in MLL datasets, negative labels usually dominate, which makes all the pseudo-labels tend to concentrate at a small value. To cope with this issue, we further perform column-wise min-max normalization on $\boldsymbol{Z}_t$ to obtain a balanced output.

Model Updating Step

Given an instance vector $\boldsymbol{x}$, we first feed it into a deep neural model $\hat{\boldsymbol{y}} = f(\boldsymbol{x}; \theta)$. Then, we update the model parameters to fit the pseudo-labels,

$$\min_{\theta} \mathcal{L}_{\text{Deep}}(\hat{Y}, \boldsymbol{Z}) = \frac{1}{n} \sum_{i=1}^{n} l(\hat{\boldsymbol{y}}_i, \boldsymbol{z}_i), \quad (4)$$

where $l(\cdot, \cdot)$ is the loss function. Recall that in Eq. (2), we enforce $\boldsymbol{Z}$ to be close to $\hat{\boldsymbol{Y}}$ using the Frobenius norm. Hence, a natural choice of loss is to reuse the mean squared error (MSE), which is adopted by most graph-based PML algorithms [Wang *et al.*, 2019; Xu *et al.*, 2020; Zhang and Fang, 2021]. However, in PML tasks, the disambiguated labels can still be noisy and the MSE loss leads to heavy overfitting. This is largely overlooked by previous works.

To tackle this problem, we shall have a closer look at the risk function. While MSE penalize too much on prediction mistakes, other robust surrogate losses can be considered, such as mean absolute error (MAE). However, we notice that MAE violates our smooth prediction assumption. To this end, we adopt the widely-used binary cross-entropy loss (BCE) for multi-label learning,

$$l(\hat{\boldsymbol{y}}_i, \boldsymbol{z}_i) = \sum_{j=1}^{L} -z_{ij} \log(\sigma(\hat{y}_{ij})) - (1-z_{ij}) \log(1-\sigma(\hat{y}_{ij})),$$

where $\sigma(\cdot)$ is the sigmoid function. On one hand, BCE loss approximately grows linearly for largely negative values, which makes it less sensitive to outliers. On the other hand, BCE outputs relatively smooth output such that the smoothness is preserved.

In effect, we can go further on the risk. A plethora of works [Ghosh *et al.*, 2017; Feng *et al.*, 2020] has demonstrated that

---

Algorithm 1 The pseudo-code of PLAIN
___
Output: The deep MLL network parameter θ
Input: Training dataset $\mathcal{D}$, instance graph $\boldsymbol{A}^x$, label graph $\boldsymbol{A}^y$, hyperparameters γ, α, β and η, iteration number T
1: $\boldsymbol{Z} \leftarrow \boldsymbol{Y}$
2: $\boldsymbol{L}_x \leftarrow \boldsymbol{I} - \boldsymbol{D}_x^{-\frac{1}{2}} \boldsymbol{A}^x \boldsymbol{D}_x^{-\frac{1}{2}}$
3: $\boldsymbol{L}_y \leftarrow \boldsymbol{I} - \boldsymbol{D}_y^{-\frac{1}{2}} \boldsymbol{A}^y \boldsymbol{D}_y^{-\frac{1}{2}}$
4: **for** epoch $t = 1, 2, 3, \ldots$ **do**
5:     **for** batch $j = 1, 2, \ldots$ in $\mathcal{D}$ **do**
6:         Sample a mini-batch data $(\boldsymbol{X}^j, \boldsymbol{Z}^j)$
7:         $\hat{\boldsymbol{Y}}^j \leftarrow f(\boldsymbol{X}^j; \theta)$
8:         $\theta \leftarrow \text{OPTIMIZE}(l(\hat{\boldsymbol{Y}}^j, \boldsymbol{Z}^j; \theta))$
9:     **end for**
10:   $\hat{\boldsymbol{Y}} \leftarrow f(\boldsymbol{X}; \theta)$
11:   **for** iter $i = 1, 2, \ldots, T$ **do**
12:     $\boldsymbol{Z} \leftarrow \boldsymbol{Z} - \gamma \frac{\partial \mathcal{L}_{\text{LP}}(\boldsymbol{Z})}{\partial \boldsymbol{Z}}$
13:   **end for**
14: **end for**

BCE is still less robust since it is unbounded. Let us consider the probabilistic output $\sigma(\hat{\boldsymbol{y}}_i)$ of the deep model. Then the BCE loss is equivalent to a negative log loss on $\sigma(\hat{\boldsymbol{y}}_i)$. If the predicted probability of the label is close to zero, the loss can grow quickly to infinity. Previous works have suggested that bounded loss on the probabilistic output is theoretically more robust than unbounded ones. One common choice of bounded loss is the probabilistic MSE loss (PMSE), i.e. $l(\hat{\boldsymbol{y}}_i, \boldsymbol{z}_i) = ||\sigma(\hat{\boldsymbol{y}}_i) - \boldsymbol{z}_i||_2^2$. In that case, the training criterion is consistent with our propagation error. Empirically, both PMSE and BCE loss perform well, and thus, we instantiate our model using BCE loss when comparing it with other methods. In Section 4.4, we present a comparison between these losses, which clearly verifies our conjectures.

Fast Implementation. To efficiently build the instance graph from training data, we employ a quick similarity search python package Faiss [Johnson *et al.*, 2017] to find the nearest neighbors. The sparse matrix multiplication $\boldsymbol{L}_x \boldsymbol{Z}$ can be efficiently implemented by the Scipy package. In Appendix A, we provide a detailed analysis of the complexity and convergence properties to show the efficacy of our method.

4 Experiments

In this section, we present the main empirical results of PLAIN compared with state-of-the-art baselines. More experimental results are shown in Appendix B.

4.1 Datasets

We employ a total of twelve synthetic as well as real-world datasets for comparative studies. Specifically, the synthetic ones are generated from the widely-used multi-label learning datasets. A total of nine benchmark multi-label datasets[2] are used for synthetic PML datasets generation, including Emotions, Birds, Medical, Image, Bibtex, Corel5K, Eurlex-dc,

---
[2]http://mulan.sourceforge.net/datasets-mlc.html

Datasets	PLAIN	PML-NI	HALE	ML-KNN	GLOCAL	PML-lc	PARTICLE
	Ranking Loss (the lower the better)						
Music-emotion	**0.217±0.006**	0.244±0.007	0.299±0.008	0.365±0.010	0.344±0.025	0.310±0.010	0.261±0.007
Music-style	**0.136±0.005**	0.138±0.009	0.271±0.015	0.229±0.010	0.253±0.007	0.269±0.016	0.351±0.014
Mirflickr	**0.088±0.006**	0.123±0.004	0.143±0.009	0.178±0.013	0.212±0.003	0.206±0.008	0.213±0.011
	Average Precision (the higher the better)						
Music-emotion	**0.664±0.007**	0.610±0.011	0.543±0.005	0.505±0.010	0.513±0.010	0.539±0.013	0.626±0.011
Music-style	**0.742±0.012**	0.740±0.013	0.687±0.013	0.659±0.014	0.645±0.007	0.598±0.009	0.621±0.016
Mirflickr	**0.847±0.012**	0.792±0.004	0.718±0.005	0.698±0.013	0.679±0.007	0.675±0.010	0.690±0.012
	Hamming Loss (the lower the better)						
Music-emotion	**0.191±0.004**	0.254±0.013	0.219±0.003	0.364±0.011	0.219±0.001	0.221±0.008	0.206±0.003
Music-style	**0.112±0.003**	0.157±0.010	0.122±0.005	0.844±0.010	0.144±0.001	0.162±0.006	0.137±0.007
Mirflickr	0.162±0.003	0.223±0.005	**0.156±0.003**	0.217±0.006	0.253±0.001	0.193±0.006	0.179±0.008

Table 1: Comparison of PLAIN with baselines on real-world datasets, where the best results are shown in bold face.

Metrics	Emotions	Birds	Medical	Image	Corel5k	Bibtext	Eurlex-dc	Eurlex-sm	NUS-WIDE	Sum
Ranking Loss	18/0/0	15/0/3	15/0/3	15/1/2	18/0/0	18/0/0	15/0/3	14/4/0	16/0/2	144/5/13
Average Precision	18/0/0	16/1/1	15/0/3	18/0/0	16/0/2	16/1/1	18/0/0	18/0/0	18/0/0	153/2/7
Hamming Loss	15/1/2	14/1/3	11/4/3	11/1/6	5/13/0	11/7/0	8/10/0	14/4/0	14/4/0	103/45/14
Sum	51/1/2	45/2/7	41/4/9	44/2/8	39/13/2	45/8/1	41/10/3	46/8/0	48/4/2	400/52/34

Table 2: Win/tie/loss counts of PLAIN's performance against baselines on synthetic data sets (pairwise t-test at 0.05 significance level).

Eurlex-sm, NUS-WIDE. Note that NUS-WIDE is a large-scale dataset, and we preprocess it as in [Lyu *et al.*, 2020]. For each data example in MLL datasets, we randomly select r irrelevant labels and aggregate them with the ground-truth labels to obtain a candidate set. For example, given an instance $\boldsymbol{x}_i$ and its ground-truth label set $\tilde{S}_i$, we select r labels from its complementary set $\mathcal{Y} - \tilde{S}_i$. If there are less than r irrelevant labels, i.e. $|\mathcal{Y} - \tilde{S}_i| < r$, we simply set the whole training set $\mathcal{Y}$ as the candidate set for $\boldsymbol{x}_i$. Following the experimental settings in [Lyu *et al.*, 2020], we choose $r \in \{1, 2, 3\}$, resulting in 27 (3×9) synthetic datasets. Furthermore, we conducted experiments on three real-world PML datasets [Zhang and Fang, 2021], including Music-emotion, Music-style, Mirflickr. These three datasets are derived from the image retrieval task [Huiskes and Lew, 2008], where the candidate labels are collected from web users and then further verified by human annotators to determine the ground-truth labels.

4.2 Baselines and Implementation Details

We choose six benchmark methods for comparative studies, including two MLL methods ML-KNN [Zhang and Zhou, 2007], GLOCAL [Zhu *et al.*, 2018], and four state-of-the-art PML methods PML-NI [Xie and Huang, 2022], PARTICLE [Zhang and Fang, 2021], HALE [Lyu *et al.*, 2020], and PML-lc [Xie and Huang, 2018]. In particular, PARTICLE and HALE are two advanced graph-based PML methods. PARTICLE [Zhang and Fang, 2021] adopts an instance-level label propagation algorithm to disambiguate the candidate labels. HALE [Lyu *et al.*, 2020] regards the denoising procedure as

a graph-matching procedure.

The parameter setups for the used methods are as follows. We fix $k = 10$ for all kNN-based methods, including PLAIN. For the baselines, we fix or fine-tune the hyperparameters as the suggested configurations in respective papers. For our PLAIN method, the deep model is comprised of three fully-connected layers. The hidden sizes are set as $[64, 64]$ for those datasets with less than 64 labels, and $[256, 256]$ for those datasets with more than 64 and less than 256 labels. For the remaining datasets, we set the hidden sizes as $[512, 512]$. The trade-off parameters α and β are hand-tuned from $\{0.001, 0.01, 0.1\}$. η is selected from $\{0.1, 1, 10\}$. Following [Iscen *et al.*, 2017], we set $\rho = 3$. We train our deep model via stochastic gradient descent and empirically set the learning rate as 0.01 for both propagation and deep model training procedures. The number of maximum iterations is set as $T = 200$ for small-scale datasets and $T = 50$ for the large-scale NUS-WIDE dataset. Besides, weight decay is applied with a rate of $5e^{-5}$ to avoid overfitting.

For performance measure, we use three widely-used multi-label metrics, *ranking loss*, *average precision* and *hamming loss* [Zhang and Zhou, 2014]. Finally, we perform ten-fold cross-validation on each dataset and the mean metric values as well as the standard deviations are reported.

4.3 Experimental Results

Table 1 reports the experimental results on real-world datasets. Table 3 lists the results on synthetic datasets, where the parameter is configured with $r = 3$. The results of $r = 1, 2$ are consistent with those of $r = 3$ and thus are

Datasets	PLAIN	PML-NI	HALE	ML-KNN	GLOCAL	PML-lc	PARTICLE
	Ranking Loss (the lower the better)						
Emotions	**0.158±0.024**	0.214±0.029	0.235±0.037	0.241±0.026	0.322±0.053	0.459±0.035	0.259±0.019
Birds	0.205±0.036	**0.187±0.035**	0.271±0.061	0.304±0.048	0.302±0.041	0.321±0.021	0.301±0.032
Medical	0.050±0.011	**0.039±0.013**	0.169±0.025	0.088±0.019	0.068±0.005	0.056±0.012	0.100±0.021
Image	**0.190±0.015**	0.289±0.018	0.192±0.015	0.342±0.026	0.264±0.010	0.467±0.025	0.315±0.073
Corel5K	**0.120±0.004**	0.205±0.006	0.259±0.011	0.139±0.007	0.164±0.003	0.169±0.013	0.333±0.050
Bibtext	**0.079±0.005**	0.126±0.010	0.601±0.009	0.232±0.006	0.131±0.006	0.342±0.005	0.287±0.010
Eurlex-dc	0.033±0.002	**0.030±0.003**	0.079±0.002	0.086±0.012	0.150±0.003	0.071±0.013	0.061±0.023
Eurlex-sm	**0.027±0.001**	0.028±0.002	0.040±0.002	0.043±0.003	0.071±0.003	0.082±0.013	0.053±0.002
NUS-WIDE	**0.211±0.003**	0.221±0.002	0.239±0.012	0.301±0.011	0.311±0.020	-	0.240±0.015
	Average Precision (the higher the better)						
Emotions	**0.812±0.043**	0.754±0.016	0.751±0.035	0.741±0.029	0.647±0.030	0.573±0.025	0.745±0.024
Birds	**0.606±0.037**	0.580±0.062	0.505±0.054	0.453±0.052	0.390±0.060	0.388±0.035	0.431±0.051
Medical	0.781±0.039	**0.849±0.035**	0.769±0.023	0.672±0.039	0.751±0.009	0.713±0.012	0.720±0.044
Image	**0.774±0.020**	0.668±0.018	0.762±0.016	0.627±0.022	0.692±0.008	0.523±0.019	0.689±0.096
Corel5K	**0.319±0.010**	0.289±0.011	0.231±0.008	0.251±0.007	0.284±0.007	0.247±0.012	0.135±0.041
Bibtext	**0.550±0.014**	0.537±0.012	0.353±0.010	0.306±0.006	0.438±0.011	0.283±0.010	0.313±0.012
Eurlex-dc	**0.740±0.008**	0.700±0.011	0.673±0.006	0.603±0.012	0.278±0.005	0.602±0.012	0.630±0.016
Eurlex-sm	**0.802±0.007**	0.718±0.011	0.739±0.013	0.761±0.012	0.556±0.008	0.582±0.013	0.695±0.012
NUS-WIDE	**0.286±0.004**	0.274±0.003	0.230±0.007	0.171±0.015	0.177±0.024	-	0.206±0.017
	Hamming Loss (the lower the better)						
Emotions	**0.223±0.033**	0.455±0.059	0.297±0.071	0.607±0.209	0.281±0.015	0.437±0.027	0.233±0.018
Birds	0.088±0.010	0.095±0.013	0.096±0.009	**0.053±0.006**	0.096±0.008	0.132±0.012	0.142±0.018
Medical	0.019±0.001	**0.014±0.002**	0.017±0.001	0.022±0.002	0.028±0.001	0.063±0.003	0.024±0.002
Image	0.207±0.009	0.273±0.028	**0.189±0.017**	0.753±0.005	0.237±0.009	0.443±0.014	0.403±0.042
Corel5K	**0.009±0.000**	0.011±0.000	0.014±0.001	0.009±0.000	0.009±0.000	0.019±0.002	**0.009±0.000**
Bibtext	**0.013±0.000**	0.015±0.000	0.016±0.001	0.014±0.000	0.015±0.001	0.021±0.001	0.017±0.000
Eurlex-dc	**0.002±0.001**	0.010±0.001	0.004±0.003	0.009±0.005	**0.002±0.000**	0.013±0.005	0.004±0.000
Eurlex-sm	**0.006±0.002**	0.011±0.001	0.008±0.005	0.012±0.001	0.011±0.000	0.019±0.002	0.010±0.001
NUS-WIDE	**0.021±0.000**	0.035±0.002	0.031±0.008	0.049±0.016	0.022±0.001	-	0.290±0.006

$\sharp$ - means over long time consumption, and thus, the experimental results are not reported.

Table 3: Comparison of PLAIN with baselines on nine synthetic PML datasets ($r = 3$), where the best results are shown in bold face.

omitted. To better understand the superiority of the proposed method, we summarize the win/tie/loss counts on all synthetic datasets in Table 2. We also report the statistical testing results in Appendix B for analyzing the relative performance amongst competing algorithms. Based on these results, we can draw the following observations:

- The proposed PLAIN method significantly outperforms all other competing approaches. For example, on the Music-emotion and Mirflicker datasets, in terms of average precision, PLAIN achieved results superior to those of the best baselines by **6.07**% and **6.94**% respectively.

- According to Table 2, out of 486 statistical comparisons, six competing algorithms, and three evaluation metrics, PLAIN beats the competing methods in **82.30**% cases. In particular, PLAIN wins or ties the baselines in **95.68**% cases regarding average precision.

- ML-KNN and GLOCAL obtain inferior performance on many datasets, such as Medical and Eurlex-dc. The reason is that these tailored MLL models regard all the candidate labels as the ground-truth and thus tend to overfit the false-positive labels.

- Two PML methods PML-NI and HALE are strong base-

lines. In particular, PML-NI achieved the second-best results on most datasets. However, PLAIN performs significantly better than them, since their performance is restricted to the linear classifier.

- On large-scale dataset NUS-WIDE, PLAIN still outperforms other methods, which validates the effectiveness of PLAIN in handling high-volume of data.

4.4 Further Analysis

Complexity Analysis. In Table 4, we report the running time of different training stages of PLAIN, including the graph building step, label propagation step (one epoch) and the deep model training step (one epoch), on datasets with various scales. First of all, we can see that the instance graph can be fast built in a few microseconds on all datasets, with the help of Faiss package [Johnson *et al.*, 2017]. Moreover, by utilizing the sparse property of instance graph and the fast computation ability of GPUs, the time complexity of label propagation is at the same magnitude as that of deep model training. In particular, on those datasets with fewer labels, such as Music-emotion and Image, the propagation phase is as fast as the deep model traveling the whole training dataset. In Figure 2, (a) and (b) show the parameter sensitivities of
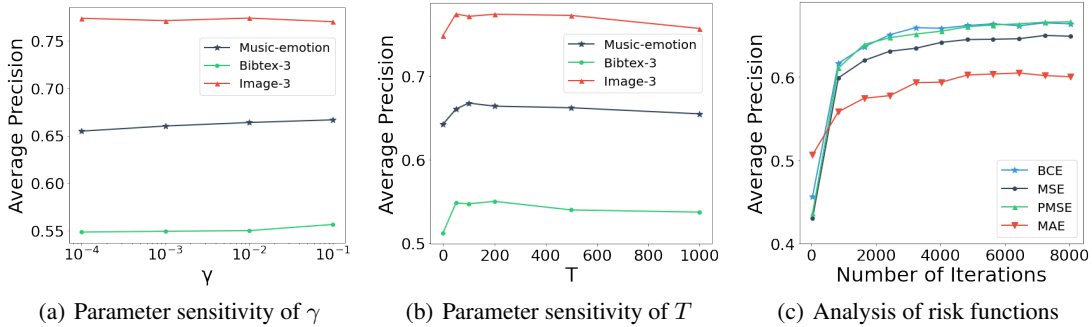
(a) Parameter sensitivity of γ (b) Parameter sensitivity of T (c) Analysis of risk functions

Figure 2: Analysis of parameter sensitivities. (a) Changes in the performance of PLAIN as γ changes; (b) Changes in the performance of PLAIN as T changes; (c) Convergence curves of PLAIN with different risk functions on Music-emotion dataset.
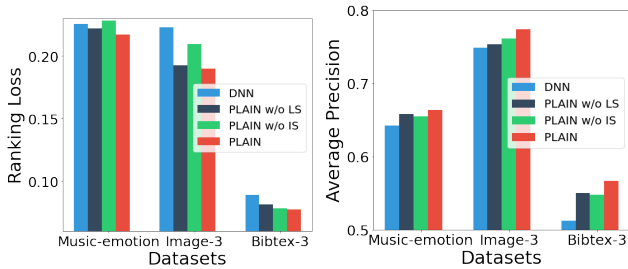

Figure 3: Ablation study on three datasets. We set $r = 3$ for Image and Bibtex. Notably, the lower ranking loss is the better and so is the hamming loss. For Average Precision, the higher the better.

Datasets	Graph Building	Label Propagation	Model Training
Image	1.0	150.2	154.4
Medical	1.1	344.0	112.2
Music-emotion	5.6	449.6	480.7
Bibtex	26.7	2498.7	600.8
Eurlex-sm	355.7	8911.2	1902.2
NUS-WIDE	1829.3	17255.2	11251.6

Table 4: Running time (**in microseconds**) of PLAIN on six datasets with various data scales, regarding the graph building, label propagation (one epoch), and deep model training (one epoch) stages.

PLAIN to T and γ. We can observe that PLAIN is stable with varying values of T and γ. Specifically, PLAIN obtains high average precision even with small T and small γ. But when $T = 0$ (without label propagation), our PLAIN model degenerates to the trivial DNN model and thus, the performance degrades. It should also be noted that when T is too large, the performance of PLAIN slightly drops because the model tends to overfit. In summary, the proposed PLAIN model requires only a few gradient steps in the label propagation phase and can be efficiently trained.

Comparison between Loss Functions. As discussed in section 3.2, mean square error (MSE) is error-prone to the label noise. Hence, we conducted experiments to compare the performance of different loss functions, including MSE, MAE, BCE and probabilistic MSE (PMSE). The results are shown in Figure 2 (c). We can see that MAE demonstrates inferior performance and low convergence. Despite its robustness, it tends to produce harsh predictions and is opposite to the graph techniques which generate smooth outputs. MSE obtains good performance, but underperforms BCE and PMSE, since it easily overfits the false-positive labels. Finally, both BCE and PMSE demonstrate good empirical results. These observations suggest that a robust loss function is essential in successful learning from PML data.

Ablation Analysis. An ablation study would be informative to see how different components of the proposed PLAIN contribute to the prediction performance. Thus, we compare

PLAIN with three variants: 1) a pure deep neural network model (DNN) that treats the candidate labels as valid ones; 2) the PLAIN model without considering label-level similarity (PLAIN w/o LS); 3) the PLAIN model without considering instance-level similarity (PLAIN w/o IS). Figure 3 reports the results on one real-world dataset as well as two synthetic datasets. In general, both graph regularizers improve the simple DNN model on most datasets and evaluation metrics. Then, by integrating the two regularizers, our PLAIN model consistently achieves the best performance. These results clearly verify the superiority of the proposed method.

5 Conclusion

In this work, we proposed a novel deep partial multi-label learning model PLAIN that copes with several limitations in existing graph-based PML methods. At each training epoch, we first involve both the instance- and label-level similarities to generate pseudo-labels via label propagation. We then train the deep model to fit the disambiguated labels. Moreover, we analyzed the time complexity and the robustness of PLAIN from both theoretical and empirical perspectives. Comprehensive experiments on synthetic datasets as well as real-world PML datasets clearly validate the efficiency, robustness, and effectiveness of our proposed method. Our work demonstrates that combining the graph and deep models together leads to high expressiveness as well as good disambiguation ability. We hope our work will increase attention toward a broader view of graph-based deep PML methods.

Acknowledgments

References

[Bai *et al.*, 2020] Junwen Bai, Shufeng Kong, and Carla P. Gomes. Disentangled variational autoencoder based multi-label classification with covariance-aware multivariate probit model. In *IJCAI*, pages 4313–4321. ijcai.org, 2020.

[Chen *et al.*, 2019a] Long Chen, Wujing Zhan, Wei Tian, Yuhang He, and Qin Zou. Deep integration: A multi-label architecture for road scene recognition. *IEEE Trans. Image Process.*, 28(10):4883–4898, 2019.

[Chen *et al.*, 2019b] Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. Multi-label image recognition with graph convolutional networks. In *CVPR*, pages 5177–5186. Computer Vision Foundation / IEEE, 2019.

[Durand *et al.*, 2019] Thibaut Durand, Nazanin Mehrasa, and Greg Mori. Learning a deep convnet for multi-label classification with partial labels. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 647–657. IEEE, 2019.

[Feng *et al.*, 2020] Lei Feng, Takuo Kaneko, Bo Han, Gang Niu, Bo An, and Masashi Sugiyama. Learning with multiple complementary labels. In *ICML*, volume 119, pages 3072–3081. PMLR, 2020.

[Ghosh *et al.*, 2017] Aritra Ghosh, Himanshu Kumar, and P. S. Sastry. Robust loss functions under label noise for deep neural networks. In *AAAI*, pages 1919–1925. AAAI Press, 2017.

[Gong *et al.*, 2022] Xiuwen Gong, Dong Yuan, and Wei Bao. Partial multi-label learning via large margin nearest neighbour embeddings. In *AAAI*, pages 6729–6736. AAAI Press, 2022.

[He *et al.*, 2019] Shuo He, Ke Deng, Li Li, Senlin Shu, and Li Liu. Discriminatively relabel for partial multi-label learning. In *ICDM*, pages 280–288. IEEE, 2019.

[Huiskes and Lew, 2008] Mark J. Huiskes and Michael S. Lew. The MIR flickr retrieval evaluation. In *SIGMM*, pages 39–43. ACM, 2008.

[Huynh and Elhamifar, 2020] Dat Huynh and Ehsan Elhamifar. Interactive multi-label CNN learning with partial labels. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9420–9429. IEEE, 2020.

[Iscen *et al.*, 2017] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Teddy Furon, and Ondrej Chum. Efficient diffusion on region manifolds: Recovering small objects with compact CNN representations. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 926–935. IEEE Computer Society, 2017.

[Johnson *et al.*, 2017] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017.

[Lai *et al.*, 2016] Hanjiang Lai, Pan Yan, Xiangbo Shu, Yunchao Wei, and Shuicheng Yan. Instance-aware hashing for multi-label image retrieval. *IEEE Trans. Image Process.*, 25(6):2469–2479, 2016.

[Lee *et al.*, 2018] Chung-Wei Lee, Wei Fang, Chih-Kuan Yeh, and Yu-Chiang Frank Wang. Multi-label zero-shot learning with structured knowledge graphs. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1576–1585. IEEE, 2018.

[Li and Wang, 2020] Ximing Li and Yang Wang. Recovering accurate labeling information from partially valid data for effective multi-label learning. In *IJCAI*, pages 1373–1380. ijcai.org, 2020.

[Li *et al.*, 2020] Ziwei Li, Gengyu Lyu, and Songhe Feng. Partial multi-label learning via multi-subspace representation. In *IJCAI*, pages 2612–2618. ijcai.org, 2020.

[Liu *et al.*, 2022] Weiwei Liu, Haobo Wang, Xiaobo Shen, and Ivor W. Tsang. The emerging trends of multi-label learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(11):7955–7974, 2022.

[Lyu *et al.*, 2020] Gengyu Lyu, Songhe Feng, and Yidong Li. Partial multi-label learning via probabilistic graph matching mechanism. In *KDD*. ACM, 2020.

[Niu *et al.*, 2019] Xuesong Niu, Hu Han, Shiguang Shan, and Xilin Chen. Multi-label co-regularization for semi-supervised facial action unit recognition. In *Adv. Neural Inform. Process. Syst.*, pages 907–917, 2019.

[Qiu *et al.*, 2018] Jiezhong Qiu, Jian Tang, Hao Ma, Yuxiao Dong, Kuansan Wang, and Jie Tang. Deepinf: Social influence prediction with deep learning. In *KDD*. ACM, 2018.

[Read *et al.*, 2011] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. Classifier chains for multi-label classification. *Mach. Learn.*, 85(3):333–359, 2011.

[Shen *et al.*, 2018] Xiaobo Shen, Weiwei Liu, Ivor W. Tsang, Quan-Sen Sun, and Yew-Soon Ong. Multilabel prediction via cross-view search. *IEEE Trans. Neural Networks Learn. Syst.*, 29(9):4324–4338, 2018.

[Shi *et al.*, 2020] Wanli Shi, Victor S. Sheng, Xiang Li, and Bin Gu. Semi-supervised multi-label learning from crowds via deep sequential generative model. In *KDD*, pages 1141–1149. ACM, 2020.

[Sun *et al.*, 2019] Lijuan Sun, Songhe Feng, Tao Wang, Congyan Lang, and Yi Jin. Partial multi-label learning by low-rank and sparse decomposition. In *AAAI*, pages 5016–5023. AAAI Press, 2019.

[Tan *et al.*, 2022] Anhui Tan, Jiye Liang, Wei-Zhi Wu, and Jia Zhang. Semi-supervised partial multi-label classification via consistency learning. *Pattern Recognit.*, 131:108839, 2022.

[Veit *et al.*, 2017] Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge J. Belongie. Learning from noisy large-scale datasets with minimal supervision. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6575–6583. IEEE Computer Society, 2017.

[Wang *et al.*, 2019] Haobo Wang, Weiwei Liu, Yang Zhao, Chen Zhang, Tianlei Hu, and Gang Chen. Discriminative and correlative partial multi-label learning. In *IJCAI*, pages 3691–3697. ijcai.org, 2019.

[Wei *et al.*, 2018] Tong Wei, Lan-Zhe Guo, Yu-Feng Li, and Wei Gao. Learning safe multi-label prediction for weakly labeled data. *Mach. Learn.*, 107(4):703–725, 2018.

[Wei *et al.*, 2021] Tong Wei, Jiang-Xin Shi, and Yu-Feng Li. Probabilistic label tree for streaming multi-label learning. In *KDD*, pages 1801–1811. ACM, 2021.

[Wei *et al.*, 2022] Tong Wei, Zhen Mao, Jiang-Xin Shi, Yu-Feng Li, and Min-Ling Zhang. A survey on extreme multi-label learning. *CoRR*, abs/2210.03968, 2022.

[Wu *et al.*, 2021] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Networks Learn. Syst.*, 32(1):4–24, 2021.

[Xie and Huang, 2018] Ming-Kun Xie and Sheng-Jun Huang. Partial multi-label learning. In *AAAI*, pages 4302–4309. AAAI Press, 2018.

[Xie and Huang, 2020] Ming-Kun Xie and Sheng-Jun Huang. Semi-supervised partial multi-label learning. In *ICDM*, pages 691–700. IEEE, 2020.

[Xie and Huang, 2022] Ming-Kun Xie and Sheng-Jun Huang. Partial multi-label learning with noisy label identification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(7):3676–3687, 2022.

[Xu *et al.*, 2020] Ning Xu, Yun-Peng Liu, and Xin Geng. Partial multi-label learning with label distribution. In *AAAI*, pages 6510–6517. AAAI Press, 2020.

[Xu *et al.*, 2021] Ning Xu, Yun-Peng Liu, Yan Zhang, and Xin Geng. Progressive enhancement of label distributions for partial multilabel learning. *IEEE Trans. Neural Networks Learn. Syst.*, pages 1–12, 2021.

[Xu *et al.*, 2022] Ning Xu, Congyu Qiao, Jiaqi Lv, Xin Geng, and Min-Ling Zhang. One positive label is sufficient: Single-positive multi-label learning with label enhancement. *CoRR*, abs/2206.00517, 2022.

[Yan and Guo, 2019] Yan Yan and Yuhong Guo. Adversarial partial multi-label learning. *CoRR*, abs/1909.06717, 2019.

[Yan *et al.*, 2021] Yan Yan, Shining Li, and Lei Feng. Partial multi-label learning with mutual teaching. *Knowl. Based Syst.*, 212:106624, 2021.

[Yeh *et al.*, 2017] Chih-Kuan Yeh, Wei-Chieh Wu, Wei-Jen Ko, and Yu-Chiang Frank Wang. Learning deep latent space for multi-label classification. In *AAAI*, pages 2838–2844. AAAI Press, 2017.

[Yilmaz *et al.*, 2021] Selim F. Yilmaz, E. Batuhan Kaynak, Aykut Koç, Hamdi Dibeklioğlu, and Suleyman Serdar Kozat. Multi-label sentiment analysis on 100 languages with dynamic weighting for label imbalance. *IEEE Trans. Neural Networks Learn. Syst.*, pages 1–13, 2021.

[Yu *et al.*, 2014] Hsiang-Fu Yu, Prateek Jain, Purushottam Kar, and Inderjit S. Dhillon. Large-scale multi-label learning with missing labels. In *ICML*, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 593–601. JMLR.org, 2014.

[Yu *et al.*, 2018] Guoxian Yu, Xia Chen, Carlotta Domeniconi, Jun Wang, Zhao Li, Zili Zhang, and Xindong Wu. Feature-induced partial multi-label learning. In *ICDM*, pages 1398–1403. IEEE Computer Society, 2018.

[Zhang and Fang, 2021] Min-Ling Zhang and Jun-Peng Fang. Partial multi-label learning via credible label elicitation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(10):3587–3599, 2021.

[Zhang and Zhou, 2007] Min-Ling Zhang and Zhi-Hua Zhou. ML-KNN: A lazy learning approach to multi-label learning. *Pattern Recognit.*, 40(7):2038–2048, 2007.

[Zhang and Zhou, 2014] Min-Ling Zhang and Zhi-Hua Zhou. A review on multi-label learning algorithms. *IEEE Trans. Knowl. Data Eng.*, 26(8):1819–1837, 2014.

[Zhu *et al.*, 2018] Yue Zhu, James T. Kwok, and Zhi-Hua Zhou. Multi-label learning with global and local label correlation. *IEEE Trans. Knowl. Data Eng.*, 30(6):1081–1094, 2018.