

Adversarial Partial Multi-Label Learning with Label Disambiguation

Yan Yan,^{*1} Yuhong Guo²

¹ School of Computer Science and Engineering, Northwestern Polytechnical University, China

² School of Computer Science, Carleton University, Canada
yanyan.nwpu@gmail.com, yuhong.guo@carleton.ca

Abstract

Partial multi-label learning (PML), which tackles the problem of learning multi-label prediction models from instances with overcomplete noisy annotations, has recently started gaining attention from the research community. In this paper, we propose a novel adversarial learning model, PML-GAN, under a generalized encoder-decoder framework for partial multi-label learning. The PML-GAN model uses a disambiguation network to identify irrelevant labels and uses a multi-label prediction network to map the training instances to their disambiguated label vectors, while deploying a generative adversarial network as an inverse mapping from label vectors to data samples in the input feature space. The learning of the overall model corresponds to a minimax adversarial game, which enhances the correspondence of input features with the output labels in a bi-directional mapping. Extensive experiments are conducted on both synthetic and real-world partial multi-label datasets, while the proposed model demonstrates the state-of-the-art performance.

Introduction

In partial multi-label learning (PML), each training instance is assigned multiple candidate labels which are only partially relevant; that is, some irrelevant noise labels are assigned together with the ground-truth labels. As it is typically difficult and costly to precisely annotate instances for multi-label data (Xie and Huang 2018), the task of PML naturally arises in many real-world scenarios with crowdsourcing annotations. In such a scenario, in order to collect the complete set of positive labels for each data instance, one can gather all labels provided by multiple annotators to form the candidate label set, which is usually overcomplete and contains additional noisy labels beyond all the true labels, leading to the PML problem. Figure 1 presents such example of an over-completely annotated training image for object recognition, where the candidate labels provided by crowdsourcing annotators cover all the ground truth labels (in black color) and some irrelevant noise labels (in red color). PML is much more challenging than standard multi-label learning as the true labels are hidden among irrelevant labels and the number of true labels is unknown. The goal of PML is to learn a

good multi-label prediction model from such a partial label training set, and hence reduce the annotation cost.

An intuitive strategy of PML is to treat all candidate labels as relevant ground truth, thus any off-the-shelf multi-label classification method can be adapted to induce an expected multi-label predictor (Zhang and Zhou 2014). This strategy, though simple, cannot work well since taking the noise labels as part of the true labels will mislead the multi-label training and induce inferior prediction models. The PML work in (Xie and Huang 2018) assumes that each candidate label has a confidence score of being a true label, and learns the confidence scores and the classifier in an alternative manner by minimizing a confidence weighted ranking loss. Although this work yields some reasonable results, the estimation of label confidence scores is error-prone, especially when noise labels dominate, which can seriously impair the classifier’s performance. The recent work in (Xie and Huang 2020) proposes to perform ground-truth label recovery and noise label identification simultaneously by exploring the label correlations and the relationships between the noise labels and feature representations. Another recent work in (Fang and Zhang 2019) presents a two-stage PML method. It estimates the confidence values of the candidate labels using iterative label propagation and then chooses the highly confident candidate labels as credible labels to induce a multi-label prediction model. This work however suffers from the cumulative errors induced in propagation, which can impact the label confidence estimation and consequently impair the prediction.

In this paper, we propose a novel adversarial learning model, PML-GAN, under a generalized encoder-decoder framework to tackle the partial multi-label learning problem. The PML-GAN model comprises four component networks: a disambiguation network that predicts the probability of each candidate label being an additive noise for a training instance; a prediction network that predicts the disambiguated *true* labels of each instance from its input features; a generation network that generates samples in the feature space given latent vectors in the label space; and a discrimination network that separates the generated samples from the real data. The prediction network and disambiguation network together form an encoder that maps data samples in the input feature space to the disambiguated label vectors, while the generation network and discrimination network form a

^{*}Work conducted while visiting Carleton University.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Example of an annotated image under the partial multi-label learning (PML) setting.

generative adversarial network (GAN) as an inverse decoding mapping from vectors in the multi-label space to samples in the input feature space. The learning of the overall model corresponds to a minimax adversarial game, which enhances the correspondence of input features with the output labels through the bi-directional encoder-decoder mapping mechanism, and consequently boosts multi-label prediction performance. To the best of our knowledge, this is the first work that exploits a generative adversarial model based bi-directional mapping mechanism for PML. We conduct extensive experiments on multiple multi-label datasets under partial multi-label learning setting. The empirical results show the proposed PML-GAN yields the state-of-the-art PML performance.

Related Work

Multi-label learning is a prevalent classification problem in many real world domains, where each instance can be assigned into multiple classes simultaneously. Many multi-label learning methods developed in the literature exploit label correlations at different degrees to produce multi-label classifiers (Zhang and Zhou 2014), including the first order methods (Zhang et al. 2018), second order methods (Li, Zhao, and Guo 2014), and high-order methods (Burkhardt and Kramer 2018). Nevertheless, standard multi-label learning methods all assume each training instance is annotated with a complete set of ground truth labels, which can be impractical in many domains, where the annotations are obtained through crowdsourcing. With the union of annotations produced by multiple noisy labelers under the crowdsourcing setting, the partial multi-label learning (PML) problem arises naturally in real world scenarios, where the set of labels assigned to each training instance not only contain the ground truth labels, but also some additional irrelevant labels.

PML is more challenging than standard multi-label learning. The previous PML work in (Xie and Huang 2018) proposes two methods, PML-FP and PML-LC, to estimate the label confidence values and optimize the relevance ordering of labels by exploring the structural information in both feature and label spaces. However, due to the inherent property of alternative optimization, in these methods, the estimation error of labeling confidence values can negatively impact the coupled multi-label predictor. The work in (Sun et al. 2019) denoises the observed label matrix based on low-rank

and sparse matrix decomposition. The recent work in (Xie and Huang 2020) proposes to learn the multi-label classifier and noisy label identifier by exploiting the label correlations as well as exploring the feature-induced noise model with the observed noise-corrupted label matrix. The work in (Xu, Liu, and Geng 2020) attempts to recover the label distributions by exploiting the topological information from the feature space and label correlations from the label space, and then induces a predictive model by fitting the recovered label distributions. Another work in (Chen et al. 2020) proposes to tackle multi-view PML problem using graph-based disambiguation. In another recent work (Fang and Zhang 2019), the authors propose to address PML problem using a two-stage strategy. It first estimates the label confidence value of each candidate label with iterative label propagation, and then performs multi-label learning over selected credible labels based on the confidence values by using pairwise label ranking (PARTICLE-VLS) or maximum a posteriori reasoning (PARTICLE-MAP). The work in (Wang et al. 2019) also presents a two-stage PML method that estimates the label confidence matrix in the first stage. However, in these two-stage methods, the confidence label estimation errors can consequently degrade the multi-label learning performance without correction interaction, especially when there are many noise labels.

Studies on weak learning, partial label learning, and noisy label learning have some connections with PML, but address different problems. Weak label learning tackles the problem of multi-label learning with incomplete labels (Sun, Zhang, and Zhou 2010; Wei et al. 2018), where some ground truth labels are missed out from the annotations. Partial label learning (PLL) tackles multi-class classification under the setting where for each training instance there is one ground-truth label among the given candidate label set (Cour, Sapp, and Taskar 2011; Liu and Dietterich 2012; Zhang and Yu 2015; Yu and Zhang 2016; Chen, Patel, and Chellappa 2018). PLL methods cannot be directly applied on the more challenging PML problems, as under PML one has unknown numbers of ground truth labels among the candidate label set for each training instance. Noisy label learning (NLL) tackles multi-class classification problems where some ground-truth labels are *replaced* by noise labels (Thekumparampil et al. 2018; Lee et al. 2018; Zhang and Sabuncu 2018; Han et al. 2018; Kaneko, Ushiku, and Harada 2019; Lee et al. 2019). The off-the-shelf NLL methods cannot be directly applied on the more challenging PML problems due to the difference of problem settings.

Generative adversarial networks (GANs) (Goodfellow et al. 2014), which perform minimax adversarial training over a generation network and a discrimination network, are one of the most popular generative models since its introduction. During the past years, a vast range of GAN-based adversarial learning methods have been developed to address different tasks, including semi-supervised learning (Kumar, Sattigeri, and Fletcher 2017; Lecouat et al. 2018), unsupervised learning (Jakab et al. 2018), and learning with noisy labels (Thekumparampil et al. 2018). The proposed work in this paper however is the first one that exploits generative adversarial models for partial multi-label learning.

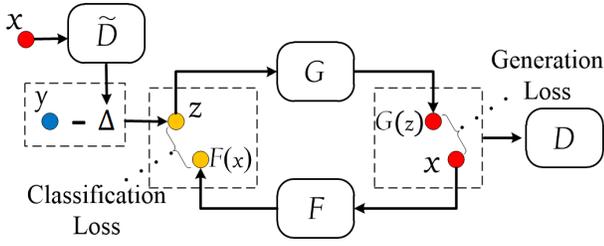


Figure 2: The proposed PML-GAN model. It has four component networks: generator G , disambiguator \tilde{D} , predictor F , and discriminator D .

Proposed Approach

In this section, we present the proposed adversarial partial multi-label learning model, PML-GAN, under the following setting. Assume we have a training set $S = (X, Y) = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, where $\mathbf{x}_i \in \mathbb{R}^d$ denotes the input feature vector for the i -th instance, and $\mathbf{y}_i \in \{0, 1\}^L$ is the corresponding annotated label indicator vector. There are multiple 1 values in each \mathbf{y}_i , which indicate either the ground truth labels or the additional mis-annotated noise labels. We aim to learn a good multi-label prediction model from this partially labeled training set.

The proposed PML-GAN model is illustrated in Figure 2, which comprises four component networks: disambiguation network \tilde{D} , prediction network F , generation network G and discrimination network D . The four components coordinate with and enhance each other under an encoder-decoder learning framework, which forms inverse mappings between the instance vectors in the input feature space and the continuous label vectors in the output class label space, aiming to facilitate the identification of a suitable prediction function. The prediction network and the disambiguation network also naturally enhance each other from a mutual learning perspective. Below we present these model components, the learning objective and the training algorithm in details.

Prediction with Disambiguated Labels

Comparing to standard multi-label learning, the main difficulty of PML is that the annotated labels $\{\mathbf{y}_i\}$ in the training data contain additive noise labels. The main challenge lies in identifying the ground truth labels \mathbf{z}_i^* from each annotated candidate label vectors \mathbf{y}_i ; that is dropping the additional 1s from each candidate label vector \mathbf{y}_i . We propose to tackle this challenge by using a disambiguation network $\tilde{D} : \Omega_{\mathbf{x}} \rightarrow \Omega_{\Delta}$ (Ω_{Δ} denotes the corresponding domain space), which predicts the irrelevant labels for a given instance. Hence the true label indicator vector \mathbf{z}_i^* can be recovered as $\mathbf{z}_i^* = \text{ReLU}(\mathbf{y}_i - \Delta_i)$ in the ideal case, where $\Delta_i \geq 0$ denotes the output of the disambiguation network $\tilde{D}(\mathbf{x}_i)$, and $\text{ReLU}(\cdot) = \max(\cdot, 0)$ denotes the commonly used rectified linear unit activation function. ReLU is used here to ensure the disambiguation effort is only counted on the candidate labels. Then we can learn a prediction network $F : \Omega_{\mathbf{x}} \rightarrow \Omega_{\mathbf{z}}$, i.e., a multi-label classifier, to predict the dis-

ambiguated ground truth labels for each instance.

Although the label indicator vectors in the training data are provided as discrete values, it is difficult for either the disambiguation network or the prediction network to directly produce discrete output values. Instead, by using a sigmoid activation function on the last layer of each network, $\tilde{D}(\mathbf{x})$ and $F(\mathbf{x})$ can predict the probability of each class label being the additive irrelevant label and the ground truth label respectively. With the disambiguation network and prediction network, we can perform partial multi-label learning by minimizing the *classification loss* on training data S :

$$\min_{F, \tilde{D}} \mathcal{L}_c(X, Y; F, \tilde{D}) = \sum_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \ell_c(F(\mathbf{x}_i), \mathbf{z}_i) \quad (1)$$

$$\text{s.t. } \mathbf{z}_i = \text{ReLU}(\mathbf{y}_i - \Delta_i), \Delta_i = \tilde{D}(\mathbf{x}_i), \forall (\mathbf{x}_i, \mathbf{y}_i) \sim S$$

where \mathbf{z}_i denotes the disambiguated label confidence vector with continuous values in $[0, 1]$, which can be viewed as a relaxation of a true label indicator vector, while $\ell_c(\cdot, \cdot)$ denotes the cross-entropy loss between the predicted probability of each label and its confidence of being a ground-truth label. We expect that the disambiguation network and the prediction network can coordinate with each other to mutually minimize this disambiguated classification loss.

Inverse Mapping with GANs

The prediction network can be viewed as an encoder that maps data samples in the input feature space to the disambiguated label vectors. To enhance the label disambiguation and hence improve multi-label classification, we propose to conduct an inverse decoding mapping from label vectors $\hat{\mathbf{z}} \in [0, 1]^L$ in the continuous label vector space to samples in the input feature space. In particular, we propose to deploy a generative adversarial network (GAN) model to transform continuous label vectors in the label space into samples in the input feature space. The GAN model comprises a generation network G and a discrimination network D . Given a label vector $\hat{\mathbf{z}}$ sampled from a prior distribution $P(\hat{\mathbf{z}})$, which can be viewed as a low-dimensional representation vector, one can generate a sample $\hat{\mathbf{x}}$ using the generation network, $\hat{\mathbf{x}} = G(\hat{\mathbf{z}})$. A two-class discriminator D is used to discriminate the generated samples from the real samples in S . The training of the GAN model is a minimax optimization problem over an *adversarial loss* function:

$$\min_G \max_D \mathcal{L}_{adv}(G, D, S) = \left[\mathbb{E}_{\mathbf{x}_i \sim S} [\log D(\mathbf{x}_i)] + \mathbb{E}_{\hat{\mathbf{z}} \sim P(\hat{\mathbf{z}})} [\log(1 - D(G(\hat{\mathbf{z}})))] \right] \quad (2)$$

where the discriminator D tries to maximally distinguish the generated samples $G(\hat{\mathbf{z}})$ from the real data samples in S , and the generator G tries to generate samples that are similar to the real data as much as possible such that the discriminator cannot tell the difference.

In theory, the samples generated by the adversarially trained generator G can have an identical distribution with the real data S (Goodfellow et al. 2014). To further ensure the generator G can provide an inverse mapping function from low-dimensional vectors in the label space to samples

Algorithm 1 Minibatch based stochastic gradient descent training algorithm for PML-GAN

Input: training set S ; trade-off parameter β ;
 k – # of update steps for the discriminator.

for # of training iterations **do**

 Sample a minibatch $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^m$ from training set S .
 Sample n label vectors $\{\hat{\mathbf{z}}_i\}_{i=1}^n$ from a prior $P(\hat{\mathbf{z}})$.

 Update the network parameters of G, \tilde{D}, F by descending with their stochastic gradients:

$$\nabla_{\Theta_{G, \tilde{D}, F}} \left\{ \begin{aligned} & \frac{1}{n} \sum_{i=1}^n \beta [\log(1 - D(G(\hat{\mathbf{z}}_i)))] + \\ & \frac{1}{m} \sum_{i=1}^m [\ell_c(F(\mathbf{x}_i), \text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))) + \\ & \ell_g(G(\text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))), \mathbf{x}_i)] \end{aligned} \right\}$$

for $r=1:k$ **do**

 Sample n label vectors $\{\hat{\mathbf{z}}_i\}_{i=1}^n$ from a prior $P(\hat{\mathbf{z}})$.
 Update the parameters of the discrimination network by ascending with its stochastic gradient:

$$\nabla_{\Theta_D} \beta \left[\frac{1}{m} \sum_{i=1}^m \log D(\mathbf{x}_i) + \frac{1}{n} \sum_{i=1}^n \log(1 - D(G(\hat{\mathbf{z}}_i))] \right]$$

end for
end for

in the feature space, we further propose to decode the disambiguated training label vectors into the training samples S with G by deploying a *generation loss*:

$$\mathcal{L}_g(G, S) = \sum_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \ell_g(G(\mathbf{z}_i), \mathbf{x}_i), \quad (3)$$

$$\text{with } \mathbf{z}_i = \text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i)), \quad (4)$$

where $\ell_g(\cdot, \cdot)$ measures the generation loss on each training instance, which can be a least squares function. This generation loss can enhance the label disambiguation and improve multi-label learning.

Learning with PML-GANs

By integrating the classification loss in Eq.(1), the adversarial loss in Eq.(2), and the generation loss in Eq.(3) together, we obtain the following minimax optimization problem for the proposed PML-GAN model:

$$\begin{aligned} & \min_{G, \tilde{D}, F} \max_D \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \left(\ell_c(F(\mathbf{x}_i), \mathbf{z}_i) + \ell_g(G(\mathbf{z}_i), \mathbf{x}_i) \right) + \\ & \beta \left(\mathbb{E}_{\mathbf{x}_i \sim S} [\log D(\mathbf{x}_i)] + \mathbb{E}_{\hat{\mathbf{z}} \sim P(\hat{\mathbf{z}})} [\log(1 - D(G(\hat{\mathbf{z}}))] \right) \quad (5) \\ & \text{s.t. } \mathbf{z}_i = \text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i)), \quad \forall (\mathbf{x}_i, \mathbf{y}_i) \sim S \end{aligned}$$

where β is a trade-off hyperparameter that controls the relative importance of the adversarial loss; the objective function can be denoted as $\mathcal{L}(G, \tilde{D}, F, D)$. The learning of the overall model corresponds to a minimax adversarial game, which enhances the bi-directional mapping between the feature and label vector spaces, and consequently boosts multi-label prediction performance.

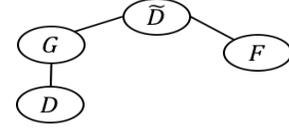


Figure 3: Dependence graph of PML-GAN.

We perform training using a minibatch based stochastic gradient descent algorithm. In each iteration of the training, the minimization over G, \tilde{D}, F and the maximization over D are conducted alternatively. The overall training algorithm is presented in Algorithm 1.

Theoretical Results

In the proposed PML-GAN model, given the generator G , the discriminator D is conditionally independent from the predictor F and the disambiguator \tilde{D} . Between G, F and \tilde{D} , G and F are conditionally independent from each other given \tilde{D} . Their independence relationship can be illustrated using the undirected dependence graph in Figure 3. Based on these conditional independence relationships, we have the following optimality results.

Proposition 1. For any G, \tilde{D} , and F , the optimal discriminator D is given by

$$D_{G, \tilde{D}, F}^*(\mathbf{x}) = D_G^*(\mathbf{x}) = p_S(\mathbf{x}) / (p_S(\mathbf{x}) + p_g(\mathbf{x})) \quad (6)$$

where $p_S(\cdot)$ and $p_g(\cdot)$ denote the probability distributions of real and generated data respectively.

Proof. Due to the conditional independence relationship between D and $\{F, \tilde{D}\}$, the optimal discriminator D only depends on the generator G . Given fixed G , the optimal discriminator can be derived in the same way as in the standard GANs (Goodfellow et al. 2014, Proposition 1). \square

Proposition 2. Assume the model has sufficient capacity. Let $C(G, \tilde{D}, F) = \max_D \mathcal{L}(G, \tilde{D}, F, D)$. Given fixed \tilde{D} , the minimum of $C(G, \tilde{D}, F)$ is lower bounded by $\mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} H(\text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))) - \beta \log 4$, which can be achieved when $F(\mathbf{x}_i) = \text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))$, $G(F(\mathbf{x}_i)) = \mathbf{x}_i$, and $p_g = p_S$.

Here $H(\cdot)$ denotes an entropy function. This proposition suggests that F and G should be inverse mapping functions for each other in the ideal optimal case.

Experiments

We conducted extensive experiments to validate the empirical performance of the proposed PML-GAN model. In this section, we report our experimental setting and results.

Experimental Setting

Datasets. We conducted experiments on twelve multi-label classification datasets. Three of them have existing partial multi-label learning settings (mirflickr, music_style and music_emotion (Fang and Zhang 2019)). For each of the

Dataset	#Inst.	#Feats	#Classes	avg.#CLs	Dataset	#Inst.	#Feats	#Classes	avg.#CLs
music_emotion	6833	98	11	5.29	music_style	6839	98	10	6.04
mirflickr	10433	100	7	3.35	image	2000	294	5	2,3,4
scene	2407	294	6	3,4,5	yeast	2417	103	14	9,10,11,12
enron	1702	1001	15	8,9,10,11,12,13	corel5k	5000	499	15	8,9,10,11,12,13
eurlex_dc	8636	100	15	8,9,10,11,12,13	eurlex_sm	12679	100	15	8,9,10,11,12,13
delicious	14000	500	15	8,9,10,11,12,13	tmc2007	28596	49060	15	8,9,10,11,12,13

Table 1: Information of the experimental data sets. The number of instances, features and classes are recorded. The “avg.#CLs” column lists the average number of candidate labels in each PML set.

other nine datasets (Zhang and Zhou 2014), we transformed it into a PML dataset by randomly adding irrelevant labels into the candidate label set of each training instance. By adding different numbers of irrelevant labels, for each dataset we can create multiple PML variants with different average numbers of candidate labels. Following the setting of (Xie and Huang 2018), we also filtered out the rare labels and kept at most 15 classes in each dataset. The detailed characteristics of the processed datasets are given in Table 1.

Comparison Methods. We compared our proposed method with five state-of-the-art PML methods and one baseline multi-label learning method. We adopted a simple but effective neural network based multi-label learning method, ML-RBF (Zhang 2009), as a baseline method, which performs PML by treating all the candidate labels as ground-truth labels. Then we used five recently developed PML methods for comparison, including the PML-LC and PML-FP methods from (Xie and Huang 2018), the PARTICLE-VLS and PARTICLE-MAP methods from (Fang and Zhang 2019), and the PML-NI from (Xie and Huang 2020).

Implementation. The proposed PML-GAN model has four component networks, all of which are designed as multilayer perceptrons with Leaky ReLU activation for the middle layers. The disambiguator, predictor, and discriminator are all three-layer networks with sigmoid activation in the output layer, while the generator is a five layer network with Tanh activation in the output layer. We used the Adam (Kingma and Ba 2014) optimizer in our implementation. The mini-batch size, m , is set to 64. The hyperparameters k (the number of steps for discriminator update) and n (the number of label vectors sampled) in Algorithm 1 are set to 1 and 2^{10} respectively. The hyperparameter β is chosen from $\{0.001, 0.01, 0.1, 1, 10\}$ based on the classification loss value \mathcal{L}_c in the training objective function; that is, the β value that leads to the smallest training \mathcal{L}_c loss will be chosen. This is a heuristic parameter selection method we create specifically for PML.

Comparison Results

We compared the proposed PML-GAN method with the six comparison methods on the twelve datasets. For each dataset, we randomly select 80% of the data for training and use the remaining 20% for testing. We repeat each experiment 10 times with different random partitions of the

datasets. The comparison test results in terms of four commonly used evaluation metrics (Hamming loss, ranking loss, one error and average precision) (Zhang and Zhou 2014) are reported in Table 2. The results are the means and standard deviations over the 10 repeated runs. We can see that the methods specially developed for PML problems all outperform the baseline multi-label neural network classifier, ML-RBF, in most cases. But it is difficult to beat the baseline competitor on all the datasets with different evaluation metrics. Among the total 48 cases over 12 datasets and 4 evaluation metrics, PML-LC, PARTICLE-VLS, PARTICLE-MAP, PML-LC and PML-FP outperform ML-RBF in 39, 42, 45, 35 and 40 cases respectively. By contrast, the proposed PML-GAN method outperforms ML-RBF *consistently* across all the 48 cases with remarkable performance gains. Even comparing with all the other five PML methods, PML-GAN produced the best results in 40 out of the total 48 cases. Moreover, the performance gains yield by PML-GAN over all the other methods are quite notable in many cases. For example, in terms of average precision, PML-GAN outperforms the best alternative comparison method by 4.6%, 4.8%, and 3.4% on *eurlex_dc*, *scene* and *image* respectively. These results clearly demonstrate the effectiveness of the proposed PML-GAN model.

The results reported in Table 2 and discussed above are produced on each dataset with a selected average number of candidate labels. As shown in Table 1, we have multiple PML variants with different numbers of candidate labels for nine of the datasets in the list. In total this provides us 49 PML datasets. We hence also conducted experiments on each of these 49 variant datasets, by comparing the proposed PML-GAN with each of the other methods in terms of the 4 evaluation metrics. In total there are 196 comparison cases for each pair of methods. For the comparison of “PML-GAN vs other method” in each case, we conducted pairwise t-test at significance level of $p < 0.05$. The win/tie/loss counts in all cases are reported in Table 3. We can see that overall the proposed PML-GAN significantly outperforms PML-NI, PARTICLE-VLS, PARTICLE-MAP, PML-LC, PML-FP, and ML-RBF in 80.6%, 75%, 77%, 81.1%, 82.6%, and 90.8% of the cases respectively. This again validates the efficacy of the proposed method for PML.

Ablation Study

As shown in Eq.(5), the objective of PML-GAN contains three parts: classification loss, generation loss and adversarial loss. The generation loss and adversarial loss are in-

Data set	avg.#C.Ls	PML-GAN	PML-NI	PARTICLE -VLS	PARTICLE -MAP	PML-LC	PML-FP	ML-RBF
Hamming loss (the smaller, the better)								
music_emotion	5.29	.200±.004	.212±.003	.212±.004	.215±.004	.236±.003	.245±.004	.779±.004
music_style	6.04	.115±.002	.116±.004	.121±.003	.175±.005	.126±.004	.126±.004	.856±.001
mirflickr	3.35	.170±.003	.167±.003	.178±.035	.189±.081	.202±.057	.202±.057	.748±.002
image	3	.202±.006	.210±.009	.234±.065	.269±.096	.264±.072	.267±.063	.754±.003
scene	4	.132±.007	.175±.003	.184±.037	.174±.035	.178±.029	.187±.038	.820±.001
yeast	10	.213±.008	.232±.004	.226±.004	.220±.008	.226±.008	.219±.009	.694±.003
enron		.186±.003	.235±.005	.197±.032	.190±.036	.206±.027	.206±.027	.813±.004
corel5k		.118±.001	.135±.003	.189±.012	.269±.027	.151±.008	.152±.008	.886±.001
eurlex_dc		.044±.001	.067±.001	.061±.001	.064±.004	.096±.001	.071±.001	.933±.001
eurlex_sm		.083±.002	.091±.008	.067±.001	.076±.002	.119±.006	.122±.002	.885±.001
delicious		.249±.002	.260±.002	.260±.003	.290±.005	.290±.004	.290±.004	.712±.002
tmc2007		.084±.001	.089±.001	.090±.003	.110±.003	.103±.002	.103±.002	.857±.001
Ranking loss (the smaller, the better)								
music_emotion	5.29	.242±.007	.251±.007	.263±.008	.240±.007	.267±.009	.275±.010	.365±.010
music_style	6.04	.145±.006	.140±.009	.163±.007	.147±.005	.215±.005	.150±.005	.242±.006
mirflickr	3.35	.124±.014	.124±.004	.227±.029	.129±.108	.160±.029	.143±.028	.195±.015
image	3	.191±.010	.217±.008	.239±.077	.250±.085	.291±.134	.217±.120	.251±.019
scene	4	.123±.009	.213±.010	.177±.049	.167±.060	.192±.032	.238±.056	.188±.014
yeast	10	.194±.011	.222±.005	.203±.007	.208±.012	.219±.011	.203±.008	.270±.007
enron		.182±.012	.236±.013	.240±.078	.182±.029	.239±.048	.239±.047	.244±.010
corel5k		.295±.011	.392±.009	.367±.032	.311±.008	.366±.035	.398±.025	.404±.082
eurlex_dc		.067±.005	.126±.010	.150±.004	.085±.004	.137±.008	.131±.001	.135±.003
eurlex_sm		.122±.007	.246±.037	.129±.007	.127±.009	.282±.007	.182±.008	.183±.003
delicious		.258±.004	.287±.002	.314±.005	.276±.004	.277±.005	.276±.005	.316±.003
tmc2007		.070±.001	.077±.001	.096±.008	.095±.007	.082±.005	.080±.005	.153±.002
One error (the smaller, the better)								
music_emotion	5.29	.450±.028	.500±.014	.473±.016	.475±.018	.556±.028	.540±.027	.587±.019
music_style	6.04	.347±.016	.355±.016	.374±.005	.399±.019	.409±.013	.408±.013	.385±.006
mirflickr	3.35	.236±.059	.307±.020	.165±.150	.229±.306	.300±.129	.298±.121	.338±.002
image	3	.342±.014	.401±.028	.369±.134	.387±.147	.542±.191	.549±.174	.398±.034
scene	4	.321±.022	.413±.018	.340±.078	.349±.082	.497±.089	.523±.118	.428±.022
yeast	10	.245±.017	.290±.009	.248±.019	.252±.018	.257±.017	.263±.027	.408±.023
enron		.307±.035	.498±.024	.411±.101	.351±.040	.494±.039	.498±.038	.495±.019
corel5k		.685±.015	.792±.016	.835±.025	.721±.035	.784±.029	.787±.024	.809±.015
eurlex_dc		.307±.013	.521±.015	.390±.016	.374±.014	.707±.014	.518±.011	.342±.008
eurlex_sm		.339±.013	.516±.019	.350±.014	.360±.015	.506±.031	.542±.018	.340±.005
delicious		.368±.009	.415±.007	.366±.015	.414±.018	.401±.015	.399±.013	.450±.009
tmc2007		.202±.007	.214±.008	.194±.029	.267±.018	.235±.019	.236±.019	.388±.006
Average precision (the larger, the better)								
music_emotion	5.29	.621±.013	.598±.007	.605±.012	.612±.009	.574±.013	.568±.014	.506±.012
music_style	6.04	.732±.010	.729±.012	.715±.009	.709±.009	.702±.008	.703±.008	.646±.010
mirflickr	3.35	.777±.027	.787±.008	.678±.027	.791±.202	.736±.043	.758±.039	.676±.048
image	3	.775±.010	.740±.013	.741±.090	.729±.086	.644±.131	.725±.119	.723±.021
scene	4	.801±.012	.688±.011	.750±.074	.753±.064	.689±.047	.710±.079	.728±.015
yeast	10	.732±.014	.701±.004	.724±.010	.714±.010	.721±.012	.728±.010	.634±.008
enron		.665±.019	.580±.009	.595±.099	.661±.047	.556±.041	.575±.041	.560±.009
corel5k		.441±.012	.345±.010	.377±.025	.415±.008	.345±.027	.384±.021	.334±.008
eurlex_dc		.797±.009	.704±.022	.692±.013	.751±.008	.693±.019	.716±.014	.710±.000
eurlex_sm		.720±.009	.558±.023	.705±.009	.683±.011	.438±.016	.679±.011	.656±.000
delicious		.630±.006	.597±.003	.596±.007	.601±.008	.607±.007	.608±.006	.576±.004
tmc2007		.821±.002	.807±.003	.799±.013	.759±.013	.793±.012	.794±.012	.662±.003

Table 2: Comparison results of in terms of Hamming loss, ranking loss, one error and average precision. The best results are presented in bold font. The average number of candidate labels is presented under the column “avg.#C.Ls”.

Evaluation Metric	PML-GAN vs -					
	PML-NI	PARTICLE-VLS	PARTICLE-MAP	PML-LC	PML-FP	ML-RBF
Hamming loss	36/11/2	39/6/4	38/9/2	40/7/2	40/3/6	45/4/0
Ranking loss	38/11/0	38/9/2	38/8/3	38/8/3	40/5/4	44/3/2
One error	44/5/0	33/12/4	39/7/3	41/8/0	40/9/0	43/6/0
Average precision	40/8/1	37/8/4	36/10/3	40/6/3	42/4/3	46/3/0
Total	158/35/3	147/35/14	151/34/11	159/29/8	162/21/13	178/16/2

Table 3: Win/tie/loss counts of pairwise t-test (with $p < 0.05$) between PML-GAN and each comparison method over all dataset variants with different numbers of candidate labels.

Data set	PML-GAN	CLS-GEN	CLS-GAN	CLS-ML	PML-GAN	CLS-GEN	CLS-GAN	CLS-ML
Hamming loss (the smaller, the better)					Ranking loss (the smaller, the better)			
music_emotion	.200±.004	.203±.003	.202±.004	.207±.004	.242±.007	.249±.007	.244±.007	.250±.010
music_style	.115±.002	.118±.003	.117±.004	.121±.001	.145±.006	.147±.007	.149±.007	.155±.004
mirflickr	.170±.003	.173±.004	.174±.005	.177±.004	.124±.014	.131±.019	.133±.021	.136±.019
image	.202±.006	.206±.005	.204±.008	.220±.006	.191±.010	.195±.010	.196±.016	.201±.010
scene	.132±.007	.140±.010	.138±.005	.148±.008	.123±.009	.130±.007	.137±.010	.140±.020
yeast	.213±.008	.219±.007	.216±.003	.222±.006	.194±.011	.199±.008	.195±.005	.203±.006
enron	.186±.003	.273±.015	.277±.014	.281±.012	.182±.012	.185±.009	.188±.009	.189±.009
corel5k	.118±.001	.118±.001	.120±.003	.122±.003	.295±.011	.304±.013	.299±.007	.306±.015
eurlex_dc	.044±.001	.050±.001	.047±.001	.054±.001	.067±.005	.068±.004	.068±.005	.071±.008
eurlex_sm	.083±.002	.085±.001	.086±.001	.088±.002	.122±.007	.125±.003	.125±.005	.127±.004
delicious	.249±.002	.251±.003	.252±.001	.255±.002	.258±.004	.261±.007	.259±.005	.269±.003
tmc2007	.084±.001	.086±.001	.086±.002	.091±.001	.070±.001	.073±.002	.072±.001	.075±.003
Average precision (the larger, the better)					One error (the smaller, the better)			
music_emotion	.621±.013	.608±.013	.612±.014	.605±.013	.450±.028	.465±.019	.469±.029	.478±.028
music_style	.732±.010	.725±.012	.726±.011	.720±.004	.347±.016	.359±.021	.356±.018	.367±.007
mirflickr	.777±.027	.765±.037	.761±.036	.754±.036	.336±.059	.384±.084	.399±.067	.417±.086
image	.775±.010	.766±.009	.766±.018	.758±.011	.342±.014	.359±.016	.359±.029	.364±.021
scene	.801±.012	.793±.009	.783±.013	.780±.021	.321±.022	.330±.017	.349±.020	.350±.027
yeast	.732±.014	.723±.012	.730±.008	.715±.009	.245±.017	.262±.025	.247±.012	.264±.018
enron	.665±.019	.658±.017	.648±.028	.645±.021	.307±.035	.328±.024	.347±.049	.350±.035
corel5k	.441±.012	.431±.015	.439±.009	.428±.017	.685±.015	.705±.023	.690±.018	.707±.020
eurlex_dc	.797±.009	.790±.009	.792±.009	.779±.015	.307±.013	.310±.015	.312±.015	.315±.021
eurlex_sm	.720±.009	.713±.005	.713±.008	.711±.005	.339±.013	.351±.009	.350±.011	.356±.010
delicious	.630±.006	.627±.009	.626±.005	.620±.002	.369±.009	.375±.015	.381±.012	.386±.003
tmc2007	.821±.002	.817±.003	.818±.004	.815±.004	.202±.007	.205±.004	.206±.007	.210±.007

Table 4: Comparison results of PML-GAN and its three ablation variants.

tegrated to assist the predictor training. To investigate and validate the contribution of the generation loss and adversarial loss, we conducted an ablation study by comparing PML-GAN with three of its ablation variants: (1) CLS-GEN, which drops the adversarial loss; (2) CLS-GAN, which drops the generation loss; and (3) CLS-ML, which only uses the classification loss by dropping both the adversarial loss and generation loss. The comparison results are reported in Table 4. We can see that comparing to the full model, all three variants produced inferior results in general. Among the three variants, both *CLS-GEN* and *CLS-GAN* outperform *CLS-ML* in most cases. This suggests that both the generation loss and the adversarial loss are critical terms for the proposed model. Moreover, even the baseline variant *CLS-ML* still produces some reasonable PML results. This sug-

gests the integration of our proposed prediction network and disambiguation network is also effective.

Conclusion

In this paper, we proposed a novel adversarial model for PML. The proposed model comprises four component networks, which form an encoder-decoder framework to improve noise label disambiguation and boost multi-label learning performance. The training problem forms a min-max adversarial optimization, which is solved using an alternative min-max procedure with minibatch stochastic gradient descent. We conducted extensive experiments on multiple PML datasets. The results show that the proposed model outperforms all the comparison methods and achieves the state-of-the-art PML performance.

Acknowledgments

This research was supported in part by the NSERC discovery grant, the Canada Research Chairs program, and the China Scholarship Council.

A Appendix

A.1 Proof for Proposition 2

Proof. This proposition suggests that F and G should be inverse mapping functions for each other in the ideal optimal case. Based on the solution for the optimal discriminator D^* in Proposition 1, we have:

$$\begin{aligned} & \mathcal{L}_{adv}(G, D^*, S) \\ &= \mathbb{E}_{\mathbf{x} \sim p_S} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\hat{\mathbf{z}} \sim P(\hat{\mathbf{z}})} [\log(1 - D_G^*(G(\hat{\mathbf{z}})))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_S} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim p_g} [\log(1 - D_G^*(\mathbf{x}))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_S} \left[\log \frac{p_S(\mathbf{x})}{p_S(\mathbf{x}) + p_g(\mathbf{x})} \right] + \mathbb{E}_{\mathbf{x} \sim p_g} \left[\log \frac{p_g(\mathbf{x})}{p_S(\mathbf{x}) + p_g(\mathbf{x})} \right] \end{aligned}$$

Hence,

$$\begin{aligned} C(G, \tilde{D}, F) &= \max_{\tilde{D}} \mathcal{L}(G, \tilde{D}, F, D) \\ &= \left\{ \begin{array}{l} \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \ell_c(F(\mathbf{x}_i), \text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))) + \\ \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \ell_g(G(\text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))), \mathbf{x}_i) + \\ \beta \left(\mathbb{E}_{\mathbf{x} \sim p_S} \left[\log \frac{p_S(\mathbf{x})}{p_S(\mathbf{x}) + p_g(\mathbf{x})} \right] + \mathbb{E}_{\mathbf{x} \sim p_g} \left[\log \frac{p_g(\mathbf{x})}{p_S(\mathbf{x}) + p_g(\mathbf{x})} \right] \right) \end{array} \right\} \end{aligned}$$

Note given fixed \tilde{D} , F is conditionally independent from G and D . Hence the minimization of $C(G, \tilde{D}, F)$ over F can be independently conducted from the minimization over G . Let $\mathbf{z}_i = \text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))$. With the cross-entropy loss function $\ell_c(\cdot, \cdot)$, we have:

$$\begin{aligned} & \min_F C(G, \tilde{D}, F) \\ &\equiv \min_F \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \ell_c(F(\mathbf{x}_i), \text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))) \\ &\equiv \min_F \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \left[-\mathbf{z}_i^\top \log F(\mathbf{x}_i) - (1 - \mathbf{z}_i)^\top \log(1 - F(\mathbf{x}_i)) \right] \\ &\equiv \min_F \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} H(\mathbf{z}_i) + \text{KL}(\mathbf{z}_i \parallel F(\mathbf{x}_i)) \\ &\geq \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} H(\mathbf{z}_i) \end{aligned} \quad (7)$$

where $H(\cdot)$ denotes the entropy over a binomial distribution vector and $\text{KL}(\cdot)$ denotes the KL-divergence between two sets of binomial distributions. Assume sufficient capacity for F , in the ideal case the minimum can be reached when the predictor obtains the same distributions as the \mathbf{z}_i ; that is

$$F^*(\mathbf{x}_i) = \mathbf{z}_i = \text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i)), \quad \forall (\mathbf{x}_i, \mathbf{y}_i) \in S \quad (8)$$

Next let's consider the minimization problem over G . Note G is involved in both the generation loss and adversarial loss. If we could find solutions that lead to minimal in both losses separately, we can guarantee a minimal in the united loss. Based on (Goodfellow et al. 2014, Theorem 1), the adversarial loss part in $C(G, \tilde{D}, F)$ can be rewritten as

$$\begin{aligned} & \beta \mathcal{L}_{adv} \\ &= \beta \left(\mathbb{E}_{\mathbf{x} \sim p_S} \left[\log \frac{p_S(\mathbf{x})}{p_S(\mathbf{x}) + p_g(\mathbf{x})} \right] + \mathbb{E}_{\mathbf{x} \sim p_g} \left[\log \frac{p_g(\mathbf{x})}{p_S(\mathbf{x}) + p_g(\mathbf{x})} \right] \right) \\ &= \beta \left(\text{KL}(p_S, \frac{p_S + p_g}{2}) - \log 2 + \text{KL}(p_g, \frac{p_S + p_g}{2}) - \log 2 \right) \\ &= \beta \left(\text{KL}(p_S, \frac{p_S + p_g}{2}) + \text{KL}(p_g, \frac{p_S + p_g}{2}) - \log 4 \right) \\ &\geq -\beta \log 4 \end{aligned} \quad (9)$$

where the minimal can be achieved when $p_S = p_g$ which leads to zero KL-divergence values. The generation loss part (with least squares loss function) in $C(G, \tilde{D}, F)$ can be rewritten as

$$\begin{aligned} & \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \left(\ell_g(G(\text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))), \mathbf{x}_i) \right) \\ &= \mathbb{E}_{(\mathbf{x}_i, \mathbf{y}_i) \sim S} \|G(\text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))) - \mathbf{x}_i\|^2 \\ &\geq 0 \end{aligned} \quad (10)$$

where the minimal 0 can only be achieved when

$$G(\text{ReLU}(\mathbf{y}_i - \tilde{D}(\mathbf{x}_i))) = \mathbf{x}_i, \quad \forall (\mathbf{x}_i, \mathbf{y}_i) \in S \quad (11)$$

The optimal condition above can be satisfied simultaneously together with the condition $p_g = p_S$. Together with the condition in (8), these conditions lead to a lower bound of $C(G, \tilde{D}, F)$ and the proposition is proved. \square

A.2 Network Architecture Information for the PML-GAN Model

The proposed PML-GAN model has four component networks, and all of them are designed as multilayer perceptrons with LeakyReLU activation function for the middle layers. The disambiguator, predictor, and discriminator are three-layer networks with sigmoid activation in the output layer, while the generator is a five layer network with Tanh activation in the output layer. Batch normalization is also deployed in the middle three layers of the generation network. The detailed input and output dimension information for each layer of the networks is given in Table 5.

Generator G			
Input	Output	BN	Act.
$\hat{\mathbf{z}}$	512	\times	LReLU
512	1024	\checkmark	LReLU
1024	256	\checkmark	LReLU
256	128	\checkmark	LReLU
128	dim	\times	Tanh
Discriminator D			
data	512	-	LReLU
512	256	-	LReLU
256	1	-	Sigmoid
Disambiguator \tilde{D}			
x	512	-	LReLU
512	256	-	LReLU
256	class num	-	Sigmoid
Predictor F			
x	512	-	LReLU
512	256	-	LReLU
256	class num	-	Sigmoid

Table 5: The network architecture of PML-GAN. BN: Batch normalization; LReLU: Leaky rectified linear unit; Act.: Activation function; dim: Feature dimension of training samples \mathbf{x} ; class num: the number of class labels.

References

- Burkhardt, S.; and Kramer, S. 2018. Online Multi-Label Dependency Topic Models for Text Classification. *Machine Learning* 107(5): 859–886.
- Chen, C.-H.; Patel, V. M.; and Chellappa, R. 2018. Learning from Ambiguously Labeled Face Images. *IEEE transactions on pattern analysis and machine intelligence (TPAMI)* 40(7): 1653–1667.
- Chen, Z.-S.; Wu, X.; Chen, Q.-G.; Hu, Y.; and Zhang, M.-L. 2020. Multi-View Partial Multi-label Learning with Graph-based Disambiguation. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Cour, T.; Sapp, B.; and Taskar, B. 2011. Learning from Partial Labels. *Journal of Machine Learning Research (JMLR)* 12(May): 1501–1536.
- Fang, J.-P.; and Zhang, M.-L. 2019. Partial Multi-Label Learning via Credible Label Elicitation. In *Proc of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In *Proc. of Advances in neural information processing systems (NeurIPS)*.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I.; and Sugiyama, M. 2018. Co-teaching: Robust Training of Deep Neural Networks with Extremely Noisy Labels. In *Proc. of Advances in Neural Information Processing Systems (NeurIPS)*.
- Jakab, T.; Gupta, A.; Bilen, H.; and Vedaldi, A. 2018. Un-supervised Learning of Object Landmarks through Conditional Image Generation. In *Proc. of Advances in Neural Information Processing Systems (NeurIPS)*.
- Kaneko, T.; Ushiku, Y.; and Harada, T. 2019. Label-Noise Robust Generative Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*.
- Kumar, A.; Sattigeri, P.; and Fletcher, P. T. 2017. Semi-supervised Learning with GANs: Manifold Invariance with Improved Inference. In *Proc. of Advances in Neural Information Processing Systems (NeurIPS)*.
- Lecouat, B.; Foo, C.-S.; Zenati, H.; and Chandrasekhar, V. R. 2018. Semi-Supervised Learning with GANs: Revisiting Manifold Regularization. *arXiv preprint arXiv:1805.08957*.
- Lee, K.; Yun, S.; Lee, K.; Lee, H.; Li, B.; and Shin, J. 2018. Robust Determinantal Generative Classifier for Noisy Labels and Adversarial Attacks. In *Proc. of Advances in Neural Information Processing Systems (NeurIPS)*.
- Lee, K.; Yun, S.; Lee, K.; Lee, H.; Li, B.; and Shin, J. 2019. Robust Inference via Generative Classifiers for Handling Noisy Labels. In *Proc. of the International Conference on Machine Learning (ICML)*.
- Li, X.; Zhao, F.-P.; and Guo, Y.-H. 2014. Multi-label Image Classification with A Probabilistic Label Enhancement Model. In *Proc of the Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Liu, L.-P.; and Dietterich, T. 2012. A Conditional Multinomial Mixture Model for Superset Label Learning. In *Proc. of Advances in neural information processing systems (NeurIPS)*.
- Sun, L.-J.; Feng, S.-H.; Wang, T.; Lang, C.-Y.; and Jin, Y. 2019. Partial Multi-Label Learning by Low-Rank and Sparse Decomposition. In *Proc of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Sun, Y.-Y.; Zhang, Y.; and Zhou, Z.-H. 2010. Multi-label Learning with Weak Label. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Thekumparampil, K.; Khetan, A.; Lin, Z.; and Oh, S. 2018. Robustness of Conditional GANs to Noisy Labels. In *Proc. of Advances in Neural Information Processing Systems (NeurIPS)*.
- Wang, H.; Liu, W.; Zhao, Y.; Zhang, C.; Hu, T.; and Chen, G. 2019. Discriminative and Correlative Partial Multi-label Learning. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Wei, T.; Guo, L.-Z.; Li, Y.-F.; and Gao, W. 2018. Learning Safe Multi-label Prediction for weakly Labeled Data. *Machine Learning* 107(4): 703–725.
- Xie, M.-K.; and Huang, S.-J. 2018. Partial Multi-Label Learning. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Xie, M.-K.; and Huang, S.-J. 2020. Partial Multi-label Learning with Noisy Label Identification. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Xu, N.; Liu, Y.-P.; and Geng, X. 2020. Partial Multi-Label Learning with Label Distribution. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Yu, F.; and Zhang, M.-L. 2016. Maximum Margin Partial Label Learning. In *Proc. of the Asian Conference on Machine Learning (ACML)*.
- Zhang, M.-L. 2009. ML-RBF: RBF Neural Networks for Multi-Label Learning. *Neural Processing Letters* 29(2): 61–74.
- Zhang, M.-L.; Li, Y.-K.; Liu, X.-Y.; and Geng, X. 2018. Binary Relevance for Multi-Label Learning: An Overview. *Frontiers of Computer Science* 12(2): 191–202.
- Zhang, M.-L.; and Yu, F. 2015. Solving the Partial Label Learning Problem: An Instance-based Approach. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Zhang, M.-L.; and Zhou, Z.-H. 2014. A Review on Multi-Label Learning Algorithms. *IEEE transactions on knowledge and data engineering* 26(8): 1819–1837.
- Zhang, Z.; and Sabuncu, M. 2018. Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels. In *Proc. of Advances in Neural Information Processing Systems (NeurIPS)*.