

---

# Binary Classification with Confidence Difference

---

Wei Wang<sup>1,2</sup>, Lei Feng<sup>3,2</sup>, Yuchen Jiang<sup>4</sup>, Gang Niu<sup>2</sup>, Min-Ling Zhang<sup>5</sup>, Masashi Sugiyama<sup>2,1</sup>

<sup>1</sup> The University of Tokyo, Chiba, Japan

<sup>2</sup> RIKEN Center for Advanced Intelligence Project, Tokyo, Japan

<sup>3</sup> Nanyang Technological University, Singapore

<sup>4</sup> Alibaba Group, Beijing, China

<sup>5</sup> Southeast University, Nanjing, China

wangw@g.ecc.u-tokyo.ac.jp lfengqaq@gmail.com

jiangyuchen.jyc@alibaba-inc.com gang.niu.ml@gmail.com

zhangml@seu.edu.cn sugi@k.u-tokyo.ac.jp

## Abstract

Recently, learning with *soft labels* has been shown to achieve better performance than learning with *hard labels* in terms of model generalization, calibration, and robustness. However, collecting pointwise labeling confidence for all training examples can be challenging and time-consuming in real-world scenarios. This paper delves into a novel weakly supervised binary classification problem called *confidence-difference (ConfDiff) classification*. Instead of pointwise labeling confidence, we are given only unlabeled data pairs with confidence difference that specifies the difference in the probabilities of being positive. We propose a risk-consistent approach to tackle this problem and show that the estimation error bound achieves the optimal convergence rate. We also introduce a risk correction approach to mitigate overfitting problems, whose consistency and convergence rate are also proven. Extensive experiments on benchmark data sets and a real-world recommender system data set validate the effectiveness of our proposed approaches in exploiting the supervision information of the confidence difference.

## 1 Introduction

Recent years have witnessed the prevalence of deep learning and its successful applications. However, the success is built on the basis of the collection of large amounts of data with unique and accurate labels. However, in many real-world scenarios, it is often difficult to satisfy such requirements. To circumvent the difficulty, various weakly supervised learning problems have been investigated accordingly, including but not limited to semi-supervised learning [1, 2, 3, 4], label-noise learning [5, 6, 7, 8, 9], positive-unlabeled learning [10, 11, 12], partial-label learning [13, 14, 15, 16, 17], unlabeled-unlabeled learning [18, 19], and similarity-based classification [20, 21, 22].

Learning with soft labels has been shown to achieve better performance than learning with hard labels in the context of supervised learning [23, 24], where each example is equipped with *pointwise labeling confidence* indicating the degree to which the labels describe the example. The advantages have been validated in many aspects, including model generalization [25, 26], calibration [27, 28], and robustness [29, 30]. For example, with the help of soft labels, knowledge distillation [25, 31] transfers knowledge from a large teacher network to a small student network. The student network can be trained more efficiently and reliably with the soft labels generated by the teacher network [32, 33, 34].

However, collecting a large number of training examples with pointwise labeling confidence may be demanding under many circumstances since it is challenging to describe the labeling confidence for each training example exactly [35, 36, 37]. Different annotators may give different values of pointwise

labeling confidence to the same example due to personal biases, and it has been demonstrated that skewed confidence values can harm classification performance [36]. Besides, giving pointwise labeling information to large-scale data sets is also expensive, laborious, and even unrealistic in many real-world scenarios [38, 39, 9]. On the contrary, leveraging supervision information of pairwise comparisons may ameliorate the biases of skewed pointwise labeling confidence and save labeling costs. Following this idea, we investigate a more practical problem setting for binary classification in this paper, where we are given *unlabeled data pairs with confidence difference* indicating the difference in the probabilities of being positive. Collecting confidence difference for training examples in pairs is much cheaper and more accessible than collecting pointwise labeling confidence for all the training examples.

Take click-through rate prediction in recommender systems [40, 41] for example. The combinations of users and their favorite/disliked items can be regarded as positive/negative data. Collecting training data takes work to distinguish between positive and negative data. Furthermore, the pointwise labeling confidence of training data may be difficult to be determined due to the extraordinarily sparse and class-imbalance problems [42]. Therefore, the collected confidence values may be biased. However, collecting the difference in the preference between a pair of candidate items for a given user is more accessible and may alleviate the biases. Section 4 will give a real-world case study of this problem. Take the disease risk estimation problem for another example. Given a person’s attributes, the goal is to predict the risk of having some disease. When asking doctors to annotate the probabilities of having the disease for patients, it takes work to determine the exact values of the probabilities. Furthermore, the probability values given by different doctors may differ due to their diverse backgrounds. On the other hand, it is much easier and less biased to estimate the relative difference in the probabilities of having the disease between two patients. Therefore, the problem of learning with confidence difference is of practical research value, but has yet to be investigated in the literature.

As a related work, Feng et al. [43] elaborated that a binary classifier can be learned from pairwise comparisons, termed Pcomp classification. For a pair of samples, they used a pairwise label of one is being more likely to be positive than the other. Since knowing the confidence difference implies knowing the pairwise label, our method requires stronger supervision than Pcomp classification. Nevertheless, we argue that, in many real-world scenarios, we may not only know one example is more likely to be positive than the other, but also know how much the difference in confidence is, as explained in the above examples. Therefore, the setting of the current paper is not so restrictive compared with that of Pcomp classification. Furthermore, our setting is more flexible than that of Pcomp classification from the viewpoint of data generation process. Pcomp classification limits the labels of pairs of training data to be in  $\{(+1, +1), (+1, -1), (-1, -1)\}$ . To cope with collected data with labels  $(-1, +1)$ , they either discard them or reverse them as  $(+1, -1)$ . On the contrary, our setting is more general and we take the examples from  $(-1, +1)$  also into consideration explicitly in the data distribution assumption. Figure 1 shows the results of a pilot experiment. Here, Strategy1 denotes Pcomp-Teacher [43] discarding examples from  $(-1, +1)$  while Strategy2 denotes Pcomp-Teacher reversing examples from  $(-1, +1)$ . We can observe that both strategies perform actually similarly. Since Strategy1 loses many data, it is intuitively expected that Strategy2 works much better, but the improvement is actually marginal. This may be because the training data distribution after reversing differs from the expected one. On the contrary, in this work we take examples from  $(-1, +1)$  into consideration directly, which is a more appropriate way to handle these data.



Figure 1: Experimental results for different strategies of handling examples from  $(-1, +1)$ .

Learning with pairwise comparisons has been investigated pervasively in the community [44, 45, 46, 47, 48, 49, 50], with applications in information retrieval [51], computer vision [52], regression [53, 54], crowdsourcing [55, 56], and graph learning [57]. It is noteworthy that there exist distinct differences between our work and previous works on learning with pairwise comparisons. Previous works have mainly tried to learn a ranking function that can rank candidate examples according to relevance or preference. In this paper, we try to learn a *pointwise binary classifier* by conducting *empirical risk minimization* under the binary classification setting.

Our contributions are summarized as follows:

- We investigate confidence-difference (ConfDiff) classification, a novel and practical weakly supervised learning problem, which can be solved via empirical risk minimization by constructing an *unbiased risk estimator*. The proposed approach can be equipped with any model, loss function, and optimizer flexibly.
- An estimation error bound is derived, showing that the proposed approach achieves the optimal parametric convergence rate. The robustness is further demonstrated by probing into the influence of an inaccurate class prior probability and noisy confidence difference.
- To mitigate overfitting issues, a risk correction approach [19] with consistency guarantee is further introduced. Extensive experimental results on benchmark data sets and a real-world recommender system data set validate the effectiveness of the proposed approaches.

## 2 Preliminaries

In this section, we discuss the background of binary classification, binary classification with soft labels, and Pcomp classification. Then, we elucidate the data generation process of ConfDiff classification.

### 2.1 Binary Classification

For binary classification, let  $\mathcal{X} = \mathbb{R}^d$  denote the  $d$ -dimensional feature space and  $\mathcal{Y} = \{+1, -1\}$  denote the label space. Let  $p(\mathbf{x}, y)$  denote the unknown joint probability density over random variables  $(\mathbf{x}, y) \in \mathcal{X} \times \mathcal{Y}$ . The task of binary classification is to learn a binary classifier  $g : \mathcal{X} \rightarrow \mathbb{R}$  which minimizes the following classification risk:

$$R(g) = \mathbb{E}_{p(\mathbf{x}, y)}[\ell(g(\mathbf{x}), y)], \quad (1)$$

where  $\ell(\cdot, \cdot)$  is a non-negative binary-class loss function, such as the 0-1 loss and logistic loss. Let  $\pi_+ = p(y = +1)$  and  $\pi_- = p(y = -1)$  denote the class prior probabilities for the positive and negative classes respectively. Furthermore, let  $p_+(\mathbf{x}) = p(\mathbf{x}|y = +1)$  and  $p_-(\mathbf{x}) = p(\mathbf{x}|y = -1)$  denote the class-conditional probability densities of positive and negative data respectively. Then the classification risk in Eq. (1) can be equivalently expressed as

$$R(g) = \pi_+ \mathbb{E}_{p_+(\mathbf{x})}[\ell(g(\mathbf{x}), +1)] + \pi_- \mathbb{E}_{p_-(\mathbf{x})}[\ell(g(\mathbf{x}), -1)]. \quad (2)$$

### 2.2 Binary Classification with Soft Labels

When the soft labels of training examples are accessible to the learning algorithm, taking advantage of them can often improve the generalization performance [23]. First, the classification risk in Eq. (1) can be equivalently expressed as

$$R(g) = \mathbb{E}_{p(\mathbf{x})}[p(y = +1|\mathbf{x})\ell(g(\mathbf{x}), +1) + p(y = -1|\mathbf{x})\ell(g(\mathbf{x}), -1)]. \quad (3)$$

Then, given training data equipped with confidence  $\{(\mathbf{x}_i, r_i)\}_{i=1}^n$  where  $r_i = p(y_i = +1|\mathbf{x}_i)$  is the *pointwise positive confidence* associated with  $\mathbf{x}_i$ , we minimize the following unbiased risk estimator to perform empirical risk minimization:

$$\hat{R}_{\text{soft}}(g) = \frac{1}{n} \sum_{i=1}^n (r_i \ell(g(\mathbf{x}_i), +1) + (1 - r_i) \ell(g(\mathbf{x}_i), -1)). \quad (4)$$

However, accurate pointwise positive confidence may be hard to be obtained in reality [36].

### 2.3 Pairwise-Comparison (Pcomp) Classification

In principle, collecting supervision information of pairwise comparisons is much easier and cheaper than pointwise supervision information [43]. In Pcomp classification [43], we are given pairs of unlabeled data where we know which one is more likely to be positive than the other. It is assumed that Pcomp data are sampled from labeled data pairs whose labels belong to  $\{(+1, -1), (+1, +1), (-1, -1)\}$ . Based on this assumption, the probability density of Pcomp data  $(\mathbf{x}, \mathbf{x}')$  is given as  $q(\mathbf{x}, \mathbf{x}')/(\pi_+^2 + \pi_-^2 + \pi_+\pi_-)$  where  $q(\mathbf{x}, \mathbf{x}') = \pi_+^2 p_+(\mathbf{x})p_+(\mathbf{x}') + \pi_-^2 p_-(\mathbf{x})p_-(\mathbf{x}') + \pi_+\pi_- p_+(\mathbf{x})p_-(\mathbf{x}')$ . Then, an unbiased risk estimator for Pcomp classification is derived as follows:

$$\hat{R}_{\text{Pcomp}}(g) = \frac{1}{n} \sum_{i=1}^n (\ell(g(\mathbf{x}_i), +1) + \ell(g(\mathbf{x}'_i), -1) - \pi_+ \ell(g(\mathbf{x}_i), -1) - \pi_- \ell(g(\mathbf{x}'_i), +1)). \quad (5)$$

In real-world scenarios, we may not only know one example is more likely to be positive than the other, but also know how much the difference in confidence is. Next, a novel weakly supervised learning setting named ConfDiff classification is introduced which can utilize such confidence difference.

## 2.4 Confidence-Difference (ConfDiff) Classification

In this subsection, the formal definition of confidence difference is given firstly. Then, we elaborate the data generation process of ConfDiff data.

**Definition 1** (Confidence Difference). *The confidence difference  $c(\mathbf{x}, \mathbf{x}')$  between an unlabeled data pair  $(\mathbf{x}, \mathbf{x}')$  is defined as*

$$c(\mathbf{x}, \mathbf{x}') = p(y' = +1|\mathbf{x}') - p(y = +1|\mathbf{x}). \quad (6)$$

As shown in the definition above, the confidence difference denotes the difference in the class posterior probabilities between the unlabeled data pair, which can measure how confident the pairwise comparison is. In ConfDiff classification, we are only given  $n$  unlabeled data pairs with confidence difference  $\mathcal{D} = \{((\mathbf{x}_i, \mathbf{x}'_i), c_i)\}_{i=1}^n$ . Here,  $c_i = c(\mathbf{x}_i, \mathbf{x}'_i)$  is the confidence difference for the unlabeled data pair  $(\mathbf{x}_i, \mathbf{x}'_i)$ . Furthermore, the unlabeled data pair  $(\mathbf{x}_i, \mathbf{x}'_i)$  is assumed to be drawn from a probability density  $p(\mathbf{x}, \mathbf{x}') = p(\mathbf{x})p(\mathbf{x}')$ . This indicates that  $\mathbf{x}_i$  and  $\mathbf{x}'_i$  are two i.i.d. instances sampled from  $p(\mathbf{x})$ . It is worth noting that the confidence difference  $c_i$  will be positive if the second instance  $\mathbf{x}'_i$  has a higher probability to be positive than the first instance  $\mathbf{x}_i$ , and will be negative otherwise. During the data collection process, the labeler can first sample two unlabeled data independently from the marginal distribution  $p(\mathbf{x})$ , then provide the confidence difference for them.

## 3 The Proposed Approach

In this section, we introduce our proposed approaches with theoretical guarantees. Besides, we show the influence of an inaccurate class prior probability and noisy confidence difference theoretically. Furthermore, we introduce a risk correction approach to improve the generalization performance.

### 3.1 Unbiased Risk Estimator

In this subsection, we show that the classification risk in Eq. (1) can be expressed with ConfDiff data in an equivalent way.

**Theorem 1.** *The classification risk  $R(g)$  in Eq. (1) can be equivalently expressed as*

$$R_{\text{CD}}(g) = \mathbb{E}_{p(\mathbf{x}, \mathbf{x}')} \left[ \frac{1}{2} (\mathcal{L}(\mathbf{x}, \mathbf{x}') + \mathcal{L}(\mathbf{x}', \mathbf{x})) \right], \quad (7)$$

where

$$\mathcal{L}(\mathbf{x}, \mathbf{x}') = (\pi_+ - c(\mathbf{x}, \mathbf{x}'))\ell(g(\mathbf{x}), +1) + (\pi_- - c(\mathbf{x}, \mathbf{x}'))\ell(g(\mathbf{x}'), -1).$$

Accordingly, we can derive an unbiased risk estimator for ConfDiff classification:

$$\widehat{R}_{\text{CD}}(g) = \frac{1}{2n} \sum_{i=1}^n (\mathcal{L}(\mathbf{x}_i, \mathbf{x}'_i) + \mathcal{L}(\mathbf{x}'_i, \mathbf{x}_i)). \quad (8)$$

**Minimum-variance risk estimator.** Actually, Eq. (8) is one of the candidates of the unbiased risk estimator. We introduce the following lemma:

**Lemma 1.** *The following expression is also an unbiased risk estimator:*

$$\frac{1}{n} \sum_{i=1}^n (\alpha \mathcal{L}(\mathbf{x}_i, \mathbf{x}'_i) + (1 - \alpha) \mathcal{L}(\mathbf{x}'_i, \mathbf{x}_i)), \quad (9)$$

where  $\alpha \in [0, 1]$  is an arbitrary weight.

Then, we introduce the following theorem:

**Theorem 2.** *The unbiased risk estimator in Eq. (8) has the minimum variance among all the candidate unbiased risk estimators in the form of Eq. (9) w.r.t.  $\alpha \in [0, 1]$ .*

Theorem 2 indicates the variance minimality of the proposed unbiased risk estimator in Eq. (8), and we adopt this risk estimator in the following sections.

### 3.2 Estimation Error Bound

In this subsection, we elaborate the convergence property of the proposed risk estimator  $\widehat{R}_{\text{CD}}(g)$  by giving an estimation error bound. Let  $\mathcal{G} = \{g : \mathcal{X} \mapsto \mathbb{R}\}$  denote the model class. It is assumed that there exists some constant  $C_g$  such that  $\sup_{g \in \mathcal{G}} \|g\|_\infty \leq C_g$  and some constant  $C_\ell$  such that  $\sup_{|z| \leq C_g} \ell(z, y) \leq C_\ell$ . We also assume that the binary loss function  $\ell(z, y)$  is Lipschitz continuous for  $z$  and  $y$  with a Lipschitz constant  $L_\ell$ . Let  $g^* = \arg \min_{g \in \mathcal{G}} R(g)$  denote the minimizer of the classification risk in Eq. (1) and  $\widehat{g}_{\text{CD}} = \arg \min_{g \in \mathcal{G}} \widehat{R}_{\text{CD}}(g)$  denote the minimizer of the unbiased risk estimator in Eq. (8). The following theorem can be derived:

**Theorem 3.** *For any  $\delta > 0$ , the following inequality holds with probability at least  $1 - \delta$ :*

$$R(\widehat{g}_{\text{CD}}) - R(g^*) \leq 8L_\ell \mathfrak{R}_n(\mathcal{G}) + 4C_\ell \sqrt{\frac{\ln 2/\delta}{2n}}, \quad (10)$$

where  $\mathfrak{R}_n(\mathcal{G})$  denotes the Rademacher complexity of  $\mathcal{G}$  for unlabeled data with size  $n$ .

From Theorem 3, we can observe that as  $n \rightarrow \infty$ ,  $R(\widehat{g}_{\text{CD}}) \rightarrow R(g^*)$  because  $\mathfrak{R}_n(\mathcal{G}) \rightarrow 0$  for all parametric models with a bounded norm, such as deep neural networks trained with weight decay [58]. Furthermore, the estimation error bound converges in  $\mathcal{O}_p(1/\sqrt{n})$ , where  $\mathcal{O}_p$  denotes the order in probability, which is the optimal parametric rate for empirical risk minimization without making additional assumptions [59].

### 3.3 Robustness of Risk Estimator

In the previous subsections, it was assumed that the class prior probability is known in advance. In addition, it was assumed that the ground-truth confidence difference of each unlabeled data pair is accessible. However, these assumptions can rarely be satisfied in real-world scenarios, since the collection of confidence difference is inevitably injected with noise. In this subsection, we theoretically analyze the influence of an inaccurate class prior probability and noisy confidence difference on the learning procedure. Later in Section 4.5, we will experimentally verify our theoretical findings.

Let  $\bar{\mathcal{D}} = \{((\mathbf{x}_i, \mathbf{x}'_i), \bar{c}_i)\}_{i=1}^n$  denote  $n$  unlabeled data pairs with noisy confidence difference, where  $\bar{c}_i$  is generated by corrupting the ground-truth confidence difference  $c_i$  with noise. Besides, let  $\bar{\pi}_+$  denote the inaccurate class prior probability accessible to the learning algorithm. Furthermore, let  $\bar{R}_{\text{CD}}(g)$  denote the empirical risk calculated based on the inaccurate class prior probability and noisy confidence difference. Let  $\bar{g}_{\text{CD}} = \arg \min_{g \in \mathcal{G}} \bar{R}_{\text{CD}}(g)$  denote the minimizer of  $\bar{R}_{\text{CD}}(g)$ . Then the following theorem gives an estimation error bound:

**Theorem 4.** *Based on the assumptions above, for any  $\delta > 0$ , the following inequality holds with probability at least  $1 - \delta$ :*

$$R(\bar{g}_{\text{CD}}) - R(g^*) \leq 16L_\ell \mathfrak{R}_n(\mathcal{G}) + 8C_\ell \sqrt{\frac{\ln 2/\delta}{2n}} + \frac{4C_\ell \sum_{i=1}^n |\bar{c}_i - c_i|}{n} + 4C_\ell |\bar{\pi}_+ - \pi_+|. \quad (11)$$

Theorem 4 indicates that the estimation error is bounded by twice the original bound in Theorem 3 with the mean absolute error of the noisy confidence difference and the inaccurate class prior probability. Furthermore, if  $\sum_{i=1}^n |\bar{c}_i - c_i|$  has a sublinear growth rate with high probability and the class prior probability is estimated consistently, the risk estimator can be even consistent. It elaborates the robustness of the proposed approach.

### 3.4 Risk Correction Approach

It is worth noting that the empirical risk in Eq. (8) may be negative due to negative terms, which is unreasonable because of the non-negative property of loss functions. This phenomenon will result in severe overfitting problems when complex models are adopted [19, 22, 43]. To circumvent this difficulty, we wrap the individual loss terms in Eq. (8) with *risk correction functions* proposed in Lu et al. [19], such as the rectified linear unit (ReLU) function  $f(z) = \max(0, z)$  and the absolute value function  $f(z) = |z|$ . In this way, the corrected risk estimator for ConfDiff classification can be

expressed as follows:

$$\begin{aligned} \tilde{R}_{\text{CD}}(g) = & \frac{1}{2n} (f(\sum_{i=1}^n (\pi_+ - c_i)\ell(g(\mathbf{x}_i), +1)) + f(\sum_{i=1}^n (\pi_- - c_i)\ell(g(\mathbf{x}'_i), -1))) \\ & + f(\sum_{i=1}^n (\pi_+ + c_i)\ell(g(\mathbf{x}'_i), +1)) + f(\sum_{i=1}^n (\pi_- + c_i)\ell(g(\mathbf{x}_i), -1))). \end{aligned} \quad (12)$$

**Theoretical analysis.** We assume that the risk correction function  $f(z)$  is Lipschitz continuous with Lipschitz constant  $L_f$ . For ease of notation, let  $\hat{A}(g) = \sum_{i=1}^n (\pi_+ - c_i)\ell(g(\mathbf{x}_i), +1)/2n$ ,  $\hat{B}(g) = \sum_{i=1}^n (\pi_- - c_i)\ell(g(\mathbf{x}'_i), -1)/2n$ ,  $\hat{C}(g) = \sum_{i=1}^n (\pi_+ + c_i)\ell(g(\mathbf{x}'_i), +1)/2n$ , and  $\hat{D}(g) = \sum_{i=1}^n (\pi_- + c_i)\ell(g(\mathbf{x}_i), -1)/2n$ . From Lemma 3 in Appendix A, the values of  $\mathbb{E}[\hat{A}(g)]$ ,  $\mathbb{E}[\hat{B}(g)]$ ,  $\mathbb{E}[\hat{C}(g)]$ , and  $\mathbb{E}[\hat{D}(g)]$  are non-negative. Therefore, we assume that there exist non-negative constants  $a, b, c$ , and  $d$  such that  $\mathbb{E}[\hat{A}(g)] \geq a$ ,  $\mathbb{E}[\hat{B}(g)] \geq b$ ,  $\mathbb{E}[\hat{C}(g)] \geq c$ , and  $\mathbb{E}[\hat{D}(g)] \geq d$ . Besides, let  $\tilde{g}_{\text{CD}} = \arg \min_{g \in \mathcal{G}} \tilde{R}_{\text{CD}}(g)$  denote the minimizer of  $\tilde{R}_{\text{CD}}(g)$ . Then, Theorem 5 is provided to elaborate the bias and consistency of  $\tilde{R}_{\text{CD}}(g)$ .

**Theorem 5.** *Based on the assumptions above, the bias of the risk estimator  $\tilde{R}_{\text{CD}}(g)$  decays exponentially as  $n \rightarrow \infty$ :*

$$0 \leq \mathbb{E}[\tilde{R}_{\text{CD}}(g)] - R(g) \leq 2(L_f + 1)C_\ell \Delta, \quad (13)$$

where  $\Delta = \exp(-2a^2n/C_\ell^2) + \exp(-2b^2n/C_\ell^2) + \exp(-2c^2n/C_\ell^2) + \exp(-2d^2n/C_\ell^2)$ . Furthermore, with probability at least  $1 - \delta$ , we have

$$|\tilde{R}_{\text{CD}}(g) - R(g)| \leq 2C_\ell L_f \sqrt{\frac{\ln 2/\delta}{2n}} + 2(L_f + 1)C_\ell \Delta. \quad (14)$$

Theorem 5 demonstrates that  $\tilde{R}_{\text{CD}}(g) \rightarrow R(g)$  in  $\mathcal{O}_p(1/\sqrt{n})$ , which means that  $\tilde{R}_{\text{CD}}(g)$  is biased yet consistent. The estimation error bound of  $\tilde{g}_{\text{CD}}$  is analyzed in Theorem 6.

**Theorem 6.** *Based on the assumptions above, for any  $\delta > 0$ , the following inequality holds with probability at least  $1 - \delta$ :*

$$R(\tilde{g}_{\text{CD}}) - R(g^*) \leq 8L_\ell \mathfrak{R}_n(\mathcal{G}) + 4C_\ell(L_f + 1)\sqrt{\frac{\ln 2/\delta}{2n}} + 4(L_f + 1)C_\ell \Delta. \quad (15)$$

Theorem 6 elucidates that as  $n \rightarrow \infty$ ,  $R(\tilde{g}_{\text{CD}}) \rightarrow R(g^*)$ , since  $\mathfrak{R}_n(\mathcal{G}) \rightarrow 0$  for all parametric models with a bounded norm [60] and  $\Delta \rightarrow 0$ . Furthermore, the estimation error bound converges in  $\mathcal{O}_p(1/\sqrt{n})$ , which is the optimal parametric rate for empirical risk minimization without additional assumptions [59].

## 4 Experiments

In this section, we verify the effectiveness of our proposed approaches experimentally.

### 4.1 Experimental Setup

We conducted experiments on benchmark data sets, including MNIST [61], Kuzushiji-MNIST [62], Fashion-MNIST [63], and CIFAR-10 [64]. In addition, four UCI data sets [65] were used, including Opendigits, USPS, Pendigits, and Letter. Since the data sets were originally designed for multi-class classification, we manually partitioned them into binary classes. For CIFAR-10, we used ResNet-34 [66] as the model architecture. For other data sets, we used a multilayer perceptron (MLP) with three hidden layers of width 300 equipped with the ReLU [67] activation function and batch normalization [68]. The logistic loss is utilized to instantiate the loss function  $\ell$ .

It is worth noting that confidence difference is given by labelers in real-world applications, while it was generated synthetically in this paper to facilitate comprehensive experimental analysis. We firstly trained a probabilistic classifier via logistic regression with ordinarily labeled data and the same neural network architecture. Then, we sampled unlabeled data in pairs at random, and generated the class posterior probabilities by inputting them into the probabilistic classifier. After that, we

Table 1: Classification accuracy (mean $\pm$ std) of each method on benchmark data sets with different class priors, where the best performance (excluding Oracle) is shown in bold.

| Class Prior   | Method            | MNIST                             | Kuzushiji                         | Fashion                           | CIFAR-10                          |
|---------------|-------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| $\pi_+ = 0.2$ | Pcomp-Unbiased    | 0.761 $\pm$ 0.017                 | 0.637 $\pm$ 0.052                 | 0.737 $\pm$ 0.050                 | 0.776 $\pm$ 0.023                 |
|               | Pcomp-ReLU        | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 |
|               | Pcomp-ABS         | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 |
|               | Pcomp-Teacher     | 0.965 $\pm$ 0.010                 | 0.871 $\pm$ 0.046                 | 0.853 $\pm$ 0.017                 | 0.836 $\pm$ 0.019                 |
|               | Oracle-Hard       | 0.990 $\pm$ 0.000                 | 0.939 $\pm$ 0.001                 | 0.979 $\pm$ 0.001                 | 0.894 $\pm$ 0.003                 |
|               | Oracle-Soft       | 0.989 $\pm$ 0.001                 | 0.939 $\pm$ 0.004                 | 0.979 $\pm$ 0.001                 | 0.893 $\pm$ 0.003                 |
|               | ConfDiff-Unbiased | 0.789 $\pm$ 0.041                 | 0.672 $\pm$ 0.053                 | 0.855 $\pm$ 0.024                 | 0.789 $\pm$ 0.025                 |
|               | ConfDiff-ReLU     | 0.968 $\pm$ 0.003                 | 0.860 $\pm$ 0.017                 | 0.964 $\pm$ 0.004                 | 0.844 $\pm$ 0.020                 |
|               | ConfDiff-ABS      | <b>0.975<math>\pm</math>0.003</b> | <b>0.898<math>\pm</math>0.003</b> | <b>0.965<math>\pm</math>0.002</b> | <b>0.862<math>\pm</math>0.015</b> |
| $\pi_+ = 0.5$ | Pcomp-Unbiased    | 0.712 $\pm$ 0.020                 | 0.578 $\pm$ 0.036                 | 0.723 $\pm$ 0.042                 | 0.703 $\pm$ 0.042                 |
|               | Pcomp-ReLU        | 0.502 $\pm$ 0.003                 | 0.502 $\pm$ 0.004                 | 0.500 $\pm$ 0.000                 | 0.602 $\pm$ 0.032                 |
|               | Pcomp-ABS         | 0.842 $\pm$ 0.012                 | 0.727 $\pm$ 0.006                 | 0.851 $\pm$ 0.012                 | 0.583 $\pm$ 0.018                 |
|               | Pcomp-Teacher     | 0.893 $\pm$ 0.014                 | 0.782 $\pm$ 0.046                 | 0.903 $\pm$ 0.016                 | 0.779 $\pm$ 0.016                 |
|               | Oracle-Hard       | 0.986 $\pm$ 0.000                 | 0.929 $\pm$ 0.002                 | 0.976 $\pm$ 0.001                 | 0.871 $\pm$ 0.003                 |
|               | Oracle-Soft       | 0.985 $\pm$ 0.001                 | 0.928 $\pm$ 0.002                 | 0.978 $\pm$ 0.001                 | 0.877 $\pm$ 0.002                 |
|               | ConfDiff-Unbiased | 0.911 $\pm$ 0.046                 | 0.712 $\pm$ 0.046                 | 0.896 $\pm$ 0.036                 | 0.720 $\pm$ 0.024                 |
|               | ConfDiff-ReLU     | 0.944 $\pm$ 0.011                 | 0.805 $\pm$ 0.015                 | 0.960 $\pm$ 0.003                 | 0.830 $\pm$ 0.007                 |
|               | ConfDiff-ABS      | <b>0.964<math>\pm</math>0.001</b> | <b>0.867<math>\pm</math>0.006</b> | <b>0.967<math>\pm</math>0.001</b> | <b>0.843<math>\pm</math>0.004</b> |
| $\pi_+ = 0.8$ | Pcomp-Unbiased    | 0.799 $\pm$ 0.005                 | 0.671 $\pm$ 0.029                 | 0.813 $\pm$ 0.029                 | 0.737 $\pm$ 0.022                 |
|               | Pcomp-ReLU        | 0.910 $\pm$ 0.031                 | 0.775 $\pm$ 0.022                 | 0.897 $\pm$ 0.023                 | 0.851 $\pm$ 0.010                 |
|               | Pcomp-ABS         | 0.854 $\pm$ 0.027                 | 0.838 $\pm$ 0.026                 | 0.921 $\pm$ 0.017                 | 0.849 $\pm$ 0.007                 |
|               | Pcomp-Teacher     | 0.943 $\pm$ 0.026                 | 0.814 $\pm$ 0.027                 | 0.936 $\pm$ 0.014                 | 0.821 $\pm$ 0.003                 |
|               | Oracle-Hard       | 0.991 $\pm$ 0.001                 | 0.942 $\pm$ 0.003                 | 0.979 $\pm$ 0.000                 | 0.897 $\pm$ 0.002                 |
|               | Oracle-Soft       | 0.990 $\pm$ 0.002                 | 0.945 $\pm$ 0.003                 | 0.980 $\pm$ 0.001                 | 0.904 $\pm$ 0.009                 |
|               | ConfDiff-Unbiased | 0.792 $\pm$ 0.017                 | 0.758 $\pm$ 0.033                 | 0.810 $\pm$ 0.035                 | 0.794 $\pm$ 0.012                 |
|               | ConfDiff-ReLU     | 0.970 $\pm$ 0.004                 | 0.886 $\pm$ 0.009                 | 0.970 $\pm$ 0.002                 | 0.851 $\pm$ 0.012                 |
|               | ConfDiff-ABS      | <b>0.983<math>\pm</math>0.002</b> | <b>0.915<math>\pm</math>0.001</b> | <b>0.975<math>\pm</math>0.002</b> | <b>0.874<math>\pm</math>0.011</b> |

generated confidence difference for each pair of sampled data according to Definition 1. To verify the effectiveness of our approaches under different class prior settings, we set  $\pi_+ \in \{0.2, 0.5, 0.8\}$  for all the data sets. Besides, we assumed that the class prior  $\pi_+$  was known for all the compared methods. We repeated the sampling-and-training procedure for five times, and the mean accuracy as well as the standard deviation were recorded.

We adopted the following variants of our proposed approaches: 1) ConfDiff-Unbiased, which denotes the method working by minimizing the unbiased risk estimator; 2) ConfDiff-ReLU, which denotes the method working by minimizing the corrected risk estimator with the ReLU function as the risk correction function; 3) ConfDiff-ABS, which denotes the method working by minimizing the corrected risk estimator with the absolute value function as the risk correction function. We compared our proposed approaches with several Pcomp methods [43], including Pcomp-Unbiased, Pcomp-ReLU, Pcomp-ABS, and Pcomp-Teacher. We also recorded the experimental results of supervised learning methods, including Oracle-Hard having access to ground-truth hard labels and Oracle-Soft having access to pointwise positive confidence. All the experiments were conducted on NVIDIA GeForce RTX 3090. The number of training epoches was set to 200 and we obtained the testing accuracy by averaging the results in the last 10 epoches. All the methods were implemented in PyTorch [69]. We used the Adam optimizer [70]. To ensure fair comparisons, We set the same hyperparameter values for all the compared approaches, where the details can be found in Appendix H.

## 4.2 Experimental Results

**Benchmark data sets.** Table 1 reports detailed experimental results for all the compared methods on four benchmark data sets. Based on Table 1, we can draw the following conclusions: a) On all the cases of benchmark data sets, our proposed ConfDiff-ABS method achieves superior performance against all of the other compared approaches significantly, which validates the effectiveness of our approach in utilizing supervision information of confidence difference; b) Pcomp-Teacher achieves

Table 2: Classification accuracy (mean $\pm$ std) of each method on UCI data sets with different class priors, where the best performance (excluding Oracle) is shown in bold.

| Class Prior   | Method            | Optdigits                         | USPS                              | Pendigits                         | Letter                            |
|---------------|-------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| $\pi_+ = 0.2$ | Pcomp-Unbiased    | 0.771 $\pm$ 0.016                 | 0.721 $\pm$ 0.046                 | 0.743 $\pm$ 0.057                 | 0.757 $\pm$ 0.028                 |
|               | Pcomp-ReLU        | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 |
|               | Pcomp-ABS         | 0.800 $\pm$ 0.001                 | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 | 0.800 $\pm$ 0.000                 |
|               | Pcomp-Teacher     | 0.901 $\pm$ 0.023                 | 0.894 $\pm$ 0.023                 | 0.928 $\pm$ 0.019                 | 0.883 $\pm$ 0.006                 |
|               | Oracle-Hard       | 0.990 $\pm$ 0.002                 | 0.984 $\pm$ 0.002                 | 0.997 $\pm$ 0.001                 | 0.978 $\pm$ 0.003                 |
|               | Oracle-Soft       | 0.990 $\pm$ 0.003                 | 0.984 $\pm$ 0.004                 | 0.998 $\pm$ 0.001                 | 0.971 $\pm$ 0.007                 |
|               | ConfDiff-Unbiased | 0.831 $\pm$ 0.078                 | 0.840 $\pm$ 0.078                 | 0.865 $\pm$ 0.079                 | 0.732 $\pm$ 0.053                 |
|               | ConfDiff-ReLU     | 0.953 $\pm$ 0.014                 | 0.957 $\pm$ 0.007                 | 0.987 $\pm$ 0.003                 | 0.929 $\pm$ 0.008                 |
|               | ConfDiff-ABS      | <b>0.963<math>\pm</math>0.009</b> | <b>0.960<math>\pm</math>0.005</b> | <b>0.988<math>\pm</math>0.002</b> | <b>0.942<math>\pm</math>0.007</b> |
| $\pi_+ = 0.5$ | Pcomp-Unbiased    | 0.651 $\pm$ 0.112                 | 0.671 $\pm$ 0.090                 | 0.748 $\pm$ 0.038                 | 0.632 $\pm$ 0.019                 |
|               | Pcomp-ReLU        | 0.630 $\pm$ 0.076                 | 0.554 $\pm$ 0.048                 | 0.514 $\pm$ 0.019                 | 0.525 $\pm$ 0.023                 |
|               | Pcomp-ABS         | 0.787 $\pm$ 0.031                 | 0.814 $\pm$ 0.018                 | 0.793 $\pm$ 0.017                 | 0.748 $\pm$ 0.031                 |
|               | Pcomp-Teacher     | 0.890 $\pm$ 0.009                 | 0.860 $\pm$ 0.012                 | 0.883 $\pm$ 0.018                 | 0.864 $\pm$ 0.024                 |
|               | Oracle-Hard       | 0.988 $\pm$ 0.003                 | 0.980 $\pm$ 0.003                 | 0.997 $\pm$ 0.001                 | 0.975 $\pm$ 0.001                 |
|               | Oracle-Soft       | 0.987 $\pm$ 0.003                 | 0.980 $\pm$ 0.003                 | 0.997 $\pm$ 0.001                 | 0.967 $\pm$ 0.006                 |
|               | ConfDiff-Unbiased | 0.917 $\pm$ 0.006                 | 0.936 $\pm$ 0.010                 | 0.945 $\pm$ 0.052                 | 0.755 $\pm$ 0.041                 |
|               | ConfDiff-ReLU     | 0.921 $\pm$ 0.011                 | 0.945 $\pm$ 0.009                 | 0.981 $\pm$ 0.004                 | 0.895 $\pm$ 0.006                 |
|               | ConfDiff-ABS      | <b>0.962<math>\pm</math>0.006</b> | <b>0.959<math>\pm</math>0.004</b> | <b>0.988<math>\pm</math>0.003</b> | <b>0.925<math>\pm</math>0.003</b> |
| $\pi_+ = 0.8$ | Pcomp-Unbiased    | 0.765 $\pm$ 0.023                 | 0.746 $\pm$ 0.012                 | 0.743 $\pm$ 0.026                 | 0.694 $\pm$ 0.031                 |
|               | Pcomp-ReLU        | 0.902 $\pm$ 0.017                 | 0.891 $\pm$ 0.024                 | 0.913 $\pm$ 0.023                 | 0.827 $\pm$ 0.025                 |
|               | Pcomp-ABS         | 0.894 $\pm$ 0.019                 | 0.879 $\pm$ 0.009                 | 0.911 $\pm$ 0.009                 | 0.870 $\pm$ 0.006                 |
|               | Pcomp-Teacher     | 0.918 $\pm$ 0.007                 | 0.933 $\pm$ 0.023                 | 0.903 $\pm$ 0.008                 | 0.872 $\pm$ 0.011                 |
|               | Oracle-Hard       | 0.987 $\pm$ 0.003                 | 0.983 $\pm$ 0.002                 | 0.997 $\pm$ 0.001                 | 0.976 $\pm$ 0.004                 |
|               | Oracle-Soft       | 0.986 $\pm$ 0.003                 | 0.985 $\pm$ 0.004                 | 0.998 $\pm$ 0.001                 | 0.965 $\pm$ 0.010                 |
|               | ConfDiff-Unbiased | 0.886 $\pm$ 0.037                 | 0.803 $\pm$ 0.042                 | 0.892 $\pm$ 0.096                 | 0.748 $\pm$ 0.015                 |
|               | ConfDiff-ReLU     | 0.949 $\pm$ 0.007                 | 0.958 $\pm$ 0.008                 | 0.986 $\pm$ 0.003                 | 0.927 $\pm$ 0.008                 |
|               | ConfDiff-ABS      | <b>0.964<math>\pm</math>0.005</b> | <b>0.964<math>\pm</math>0.003</b> | <b>0.987<math>\pm</math>0.002</b> | <b>0.945<math>\pm</math>0.007</b> |

superior performance against all of the other Pcomp approaches by a large margin. The excellent performance benefits from the effectiveness of consistency regularization for weakly supervised learning problems [1, 17, 71]; c) It is worth noting that the classification results of ConfDiff-ReLU and ConfDiff-ABS have smaller variances than ConfDiff-Unbiased. It demonstrates that the risk correction method can enhance the stability and robustness for ConfDiff classification.

**UCI data sets.** Table 2 reports detailed experimental results on four UCI data sets as well. From Table 2, we can observe that: a) On all the UCI data sets under different class prior probability settings, our proposed ConfDiff-ABS method achieves the best performance among all the compared approaches with significant superiority, which verifies the effectiveness of our proposed approaches again; b) The performance of our proposed approaches is more stable than the compared Pcomp approaches under different class prior probability settings, demonstrating the superiority of our methods in dealing with various kinds of data distributions.

### 4.3 Experiments on a Real-world Recommender System Data Set

We also conducted experiments on a recommender system data set to demonstrate our approach’s usefulness and promising applications in real-world scenarios. We used the KuaiRec [72] data set, a real-world recommender system data set collected from a well-known short-video mobile app. In this data set, user-item interactions are represented by watching ratios, i.e., the ratios of watching time to the entire length of videos. Such statistics could reveal the confidence of preference, and we regarded them as pointwise positive confidence. We generated pairwise confidence difference between pairs of items for a given user. We adopted the NCF [73] model as our backbone.

Table 3: Experimental results on the KuaiRec data set.

| Method        | HR    | NDCG  |
|---------------|-------|-------|
| BPR           | 0.464 | 0.256 |
| MRL           | 0.476 | 0.271 |
| Oracle-Hard   | 0.469 | 0.283 |
| Oracle-Soft   | 0.534 | 0.380 |
| Pcomp-Teacher | 0.179 | 0.066 |
| ConfDiff-ABS  | 0.570 | 0.372 |



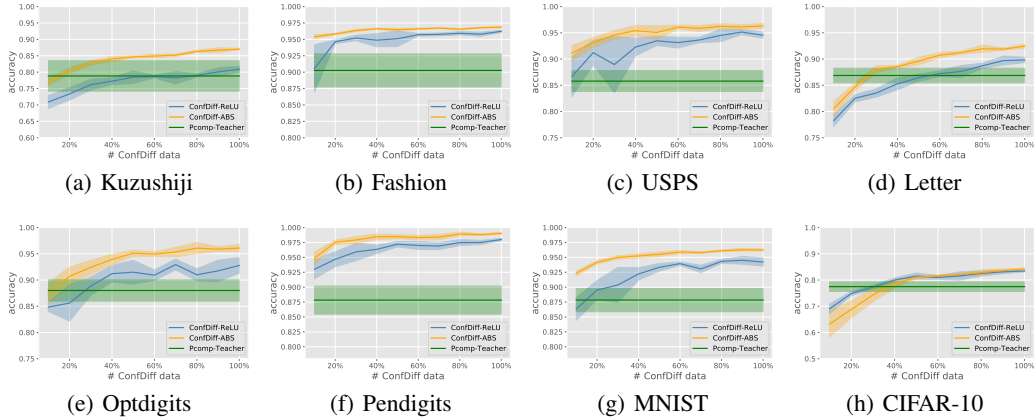


Figure 2: Classification performance of ConfDiff-ReLU and ConfDiff-ABS given a fraction of training data as well as Pcomp-Teacher given 100% of training data ( $\pi_+ = 0.5$ ).

Details of the experimental setup and the data set can be found in Appendix H. We employed 5 compared methods, including BPR [74], Margin Ranking Loss (MRL), Oracle-Hard, Oracle-Soft, and Pcomp-Teacher. Table 3 reports the hit ratio (HR) and normalized discounted cumulative gain (NDCG) results. Our approach performs comparably against Oracle in terms of NDCG and even performs better in terms of HR. On the contrary, Pcomp-Teacher does not perform well on this data set. It validates the effectiveness of our approach in exploiting the supervision information of the confidence difference.

#### 4.4 Performance with Fewer Training Data

We conducted experiments by changing the fraction of training data for ConfDiff-ReLU and ConfDiff-ABS (100% indicated that all the ConfDiff data were used for training). For comparison, we used 100% of training data for Pcomp-Teacher during the training process. Figure 2 shows the results with  $\pi_+ = 0.5$ , and more experimental results can be found in Appendix I. We can observe that the classification performance of our proposed approaches is still advantageous given a fraction of training data. Our approaches can achieve superior or comparable performance even when only 10% of training data are used. It elucidates that leveraging confidence difference may be more effective than increasing the number of training examples.

#### 4.5 Analysis on Robustness

In this subsection, we investigate the influence of an inaccurate class prior probability and noisy confidence difference on the generalization performance of the proposed approaches. Specifically, let  $\bar{\pi}_+ = \epsilon\pi_+$  denote the corrupted class prior probability with  $\epsilon$  being a real number around 1. Let  $\bar{c}_i = \epsilon'_i c_i$  denote the noisy confidence difference where  $\epsilon'_i$  is sampled from a normal distribution  $\mathcal{N}(1, \sigma^2)$ . Figure 3 shows the classification performance of our proposed approaches on MNIST and Pendigits ( $\pi_+ = 0.5$ ) with different  $\epsilon$  and  $\sigma$ . It is demonstrated that the performance degenerates with  $\epsilon = 0.8$  or  $\epsilon = 1.2$  on some data sets, which indicates that it is important to estimate the class prior accurately.

## 5 Conclusion

In this paper, we dived into a novel weakly supervised learning setting where only unlabeled data pairs equipped with confidence difference were given. To solve the problem, an unbiased risk estimator was derived to perform empirical risk minimization. An estimation error bound was established to show that the optimal parametric convergence rate can be achieved. Furthermore, a risk correction approach was introduced to alleviate overfitting issues. Extensive experimental results validated the superiority of our proposed approaches. In future, it would be promising to apply our approaches in real-world scenarios and multi-class classification settings.

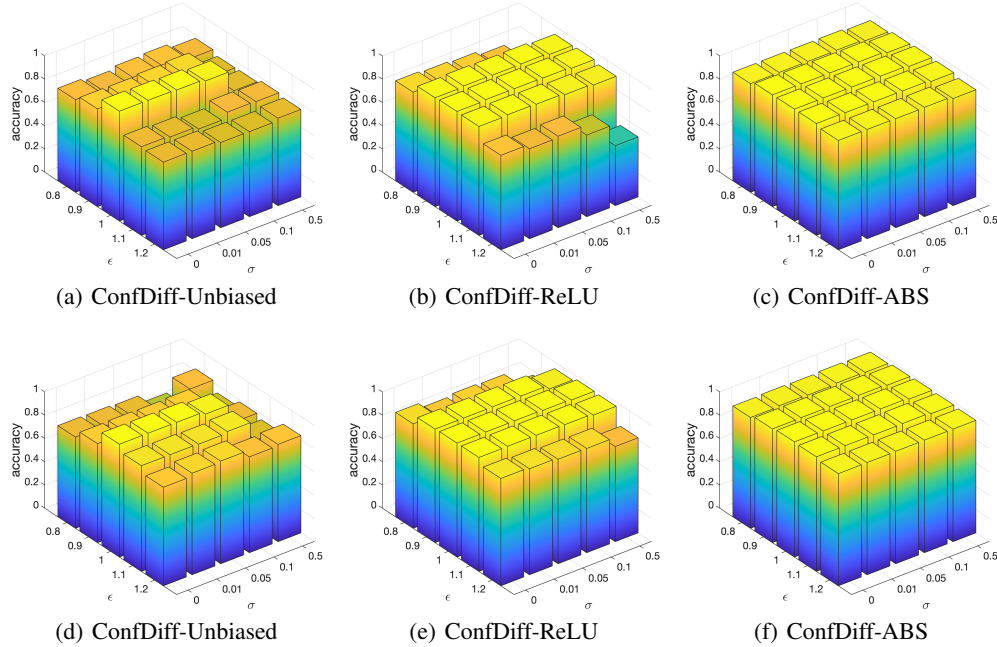


Figure 3: Classification accuracy on MNIST (the first row) and Pendigits (the second row) with  $\pi_+ = 0.5$  given an inaccurate class prior probability and noisy confidence difference.

## References

- [1] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A. Raffel. MixMatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems 32*, pages 5050–5060, 2019.
- [2] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. *Semi-Supervised Learning*. The MIT Press, 2006.
- [3] Xiaojin Zhu and Andrew B. Goldberg. Introduction to semi-supervised learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 3(1):1–130, 2009.
- [4] Yu-Feng Li and Zhi-Hua Zhou. Towards making unlabeled data never hurt. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(1):175–188, 2015.
- [5] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in Neural Information Processing Systems 31*, pages 8536–8546, 2018.
- [6] Junnan Li, Caiming Xiong, and Steven C. H. Hoi. MoPro: Webly supervised learning with momentum prototypes. In *Proceedings of the 9th International Conference on Learning Representations*, 2021.
- [7] Giorgio Patrini, Alessandro Rozza, Aditya K. Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1944–1952, 2017.
- [8] Deng-Bao Wang, Yong Wen, Lujia Pan, and Min-Ling Zhang. Learning from noisy labels with complementary loss functions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, pages 10111–10119, 2021.
- [9] Jiaheng Wei, Zhaowei Zhu, Hao Cheng, Tongliang Liu, Gang Niu, and Yang Liu. Learning with noisy labels revisited: A study using real-world human annotations. In *Proceedings of the 10th International Conference on Learning Representations*, 2022.

- [10] Marthinus C. du Plessis, Gang Niu, and Masashi Sugiyama. Analysis of learning from positive and unlabeled data. In *Advances in Neural Information Processing Systems 27*, pages 703–711, 2014.
- [11] Guangxin Su, Weitong Chen, and Miao Xu. Positive-unlabeled learning from imbalanced data. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, pages 2995–3001, 2021.
- [12] Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Gang Niu, Masashi Sugiyama, and Dacheng Tao. Rethinking class-prior estimation for positive-unlabeled learning. In *Proceedings of the 10th International Conference on Learning Representations*, 2022.
- [13] Haobo Wang, Ruixuan Xiao, Sharon Li, Lei Feng, Gang Niu, Gang Chen, and Junbo Zhao. PiCO: Contrastive label disambiguation for partial label learning. In *Proceedings of the 10th International Conference on Learning Representations*, 2022.
- [14] Timothee Cour, Ben Sapp, and Ben Taskar. Learning from partial labels. *Journal of Machine Learning Research*, 12(May):1501–1536, 2011.
- [15] Wei Wang and Min-Ling Zhang. Semi-supervised partial label learning via confidence-rated margin maximization. In *Advances in Neural Information Processing Systems 33*, pages 6982–6993, 2020.
- [16] Hongwei Wen, Jingyi Cui, Hanyuan Hang, Jiabin Liu, Yisen Wang, and Zhouchen Lin. Leveraged weighted loss for partial label learning. In *Proceedings of the 38th International Conference on Machine Learning*, pages 11091–11100, 2021.
- [17] Dong-Dong Wu, Deng-Bao Wang, and Min-Ling Zhang. Revisiting consistency regularization for deep partial label learning. In *Proceedings of the 39th International Conference on Machine Learning*, pages 24212–24225, 2022.
- [18] Nan Lu, Gang Niu, Aditya K. Menon, and Masashi Sugiyama. On the minimal supervision for training any binary classifier from only unlabeled data. In *Proceedings of the 7th International Conference on Learning Representations*, 2019.
- [19] Nan Lu, Tianyi Zhang, Gang Niu, and Masashi Sugiyama. Mitigating overfitting in supervised classification from two unlabeled datasets: A consistent risk correction approach. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pages 1115–1125, 2020.
- [20] Han Bao, Gang Niu, and Masashi Sugiyama. Classification from pairwise similarity and unlabeled data. In *Proceedings of the 35th International Conference on Machine Learning*, pages 461–470, 2018.
- [21] Han Bao, Takuya Shimada, Liyuan Xu, Issei Sato, and Masashi Sugiyama. Pairwise supervision can provably elicit a decision boundary. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, pages 2618–2640, 2022.
- [22] Yuzhou Cao, Lei Feng, Yitian Xu, Bo An, Gang Niu, and Masashi Sugiyama. Learning from similarity-confidence data. In *Proceedings of the 38th International Conference on Machine Learning*, pages 1272–1282, 2021.
- [23] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [24] Hua Yuan, Ning Xu, Yu Shi, Xin Geng, and Yong Rui. Learning from biased soft labels. In *Advances in Neural Information Processing Systems 36*, 2023, in press.
- [25] Li Yuan, Francis EH Tay, Guilin Li, Tao Wang, and Jiashi Feng. Revisiting knowledge distillation via label smoothing regularization. In *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3903–3911, 2020.

- [26] Takashi Ishida, Ikko Yamane, Nontawat Charoenphakdee, Gang Niu, and Masashi Sugiyama. Is the performance of my deep network too good to be true? A direct approach to estimating the bayes error in binary classification. In *Proceedings of the 11th International Conference on Learning Representations*, 2023.
- [27] Rafael Müller, Simon Kornblith, and Geoffrey E. Hinton. When does label smoothing help? In *Advances in Neural Information Processing Systems 32*, pages 4696–4705, 2019.
- [28] Deng-Bao Wang, Lei Feng, and Min-Ling Zhang. Rethinking calibration of deep neural networks: Do not be afraid of overconfidence. In *Advances in Neural Information Processing Systems 34*, pages 11809–11820, 2021.
- [29] Michal Lukasik, Srinadh Bhojanapalli, Aditya Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? In *Proceedings of the 37th International Conference on Machine Learning*, pages 6448–6458, 2020.
- [30] Tianyu Pang, Xiao Yang, Yinpeng Dong, Hang Su, and Jun Zhu. Bag of tricks for adversarial training. In *Proceedings of the 9th International Conference on Learning Representations*, 2021.
- [31] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [32] Mary Phuong and Christoph Lampert. Towards understanding knowledge distillation. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5142–5151, 2019.
- [33] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3967–3976, 2019.
- [34] Jianping Gou, Baosheng Yu, Stephen J. Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129:1789–1819, 2021.
- [35] Katherine M. Collins, Umang Bhatt, and Adrian Weller. Eliciting and learning with soft labels from every annotator. In *Proceedings of the 10th AAI Conference on Human Computation and Crowdsourcing*, pages 40–52, 2022.
- [36] Kazuhiko Shinoda, Hirotaka Kaji, and Masashi Sugiyama. Binary classification from positive data with skewed confidence. In *Proceedings of the 29th International Joint Conferences on Artificial Intelligence*, pages 3328–3334, 2020.
- [37] Ilya Sucholutsky, Ruairidh M. Battleday, Katherine M. Collins, Raja Marjeh, Joshua Peterson, Pulkit Singh, Umang Bhatt, Nori Jacoby, Adrian Weller, and Thomas L. Griffiths. On the informativeness of supervision signals. In *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence*, pages 2036–2046, 2023.
- [38] Davood Karimi, Haoran Dou, Simon K. Warfield, and Ali Gholipour. Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis. *Medical Image Analysis*, 65:101759, 2020.
- [39] Alexander J. Ratner, Christopher M. De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. Data programming: Creating large training sets, quickly. In *Advances in Neural Information Processing Systems 29*, pages 3567–3575, 2016.
- [40] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 52(1):1–38, 2019.
- [41] Yuchen Jiang, Qi Li, Han Zhu, Jinbei Yu, Jin Li, Ziru Xu, Huihui Dong, and Bo Zheng. Adaptive domain interest network for multi-domain recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, page 3212–3221, 2022.

- [42] Tiansheng Yao, Xinyang Yi, Derek Zhiyuan Cheng, Felix Yu, Ting Chen, Aditya Menon, Lichan Hong, Ed H. Chi, Steve Tjoa, Jieqi (Jay) Kang, and Evan Ettinger. Self-supervised learning for large-scale item recommendations. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, page 4321–4330, 2021.
- [43] Lei Feng, Senlin Shu, Nan Lu, Bo Han, Miao Xu, Gang Niu, Bo An, and Masashi Sugiyama. Pointwise binary classification with pairwise confidence comparisons. In *Proceedings of the 38th International Conference on Machine Learning*, pages 3252–3262, 2021.
- [44] Christopher J. C. Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Gregory N. Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 89–96, 2005.
- [45] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: From pairwise approach to listwise approach. In *Proceedings of the 24th International Conference on Machine Learning*, pages 129–136, 2007.
- [46] Kevin G. Jamieson and Robert D. Nowak. Active ranking using pairwise comparisons. In *Advances in Neural Information Processing Systems 24*, pages 2240–2248, 2011.
- [47] Daniel M. Kane, Shachar Lovett, Shay Moran, and Jiapeng Zhang. Active classification with comparison queries. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*, pages 355–366, 2017.
- [48] Dohyung Park, Joe Neeman, Jin Zhang, Sujay Sanghavi, and Inderjit S. Dhillon. Preference completion: Large-scale collaborative ranking from pairwise comparisons. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 1907–1916, 2015.
- [49] Nihar B. Shah, Sivaraman Balakrishnan, and Martin J. Wainwright. Feeling the bern: Adaptive estimators for bernoulli probabilities of pairwise comparisons. *IEEE Transactions on Information Theory*, 65(8):4854–4874, 2019.
- [50] Yichong Xu, Hongyang Zhang, Kyle Miller, Aarti Singh, and Artur Dubrawski. Noise-tolerant interactive learning using pairwise comparisons. In *Advances in Neural Information Processing Systems 30*, pages 2428–2437, 2017.
- [51] Tie-Yan Liu. *Learning to Rank for Information Retrieval*. Springer, 2011.
- [52] Yanwei Fu, Timothy M. Hospedales, Tao Xiang, Jiechao Xiong, Shaogang Gong, Yizhou Wang, and Yuan Yao. Robust subjective visual property prediction from crowdsourced pairwise labels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):563–577, 2015.
- [53] Liyuan Xu, Junya Honda, Gang Niu, and Masashi Sugiyama. Uncoupled regression from pairwise comparison data. In *Advances in Neural Information Processing Systems 32*, pages 3992–4002, 2019.
- [54] Yichong Xu, Sivaraman Balakrishnan, Aarti Singh, and Artur Dubrawski. Regression with comparisons: Escaping the curse of dimensionality with ordinal information. *Journal of Machine Learning Research*, 21(1):6480–6533, 2020.
- [55] Xi Chen, Paul N. Bennett, Kevyn Collins-Thompson, and Eric Horvitz. Pairwise ranking aggregation in a crowdsourced setting. In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining*, pages 193–202, 2013.
- [56] Shiwei Zeng and Jie Shen. Efficient PAC learning from the crowd with pairwise comparisons. In *Proceedings of the 39th International Conference on Machine Learning*, pages 25973–25993, 2022.
- [57] Yixuan He, Quan Gan, David Wipf, Gesine D. Reinert, Junchi Yan, and Mihai Cucuringu. GNNRank: Learning global rankings from pairwise comparisons via directed graph neural networks. In *Proceedings of the 39th International Conference on Machine Learning*, pages 8581–8612, 2022.

- [58] Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In *Proceedings of the 31st Conference On Learning Theory*, pages 297–299, 2018.
- [59] Shahar Mendelson. Lower bounds for the empirical minimization algorithm. *IEEE Transactions on Information Theory*, 54(8):3797–3803, 2008.
- [60] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. The MIT Press, 2012.
- [61] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [62] Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical Japanese literature. *arXiv preprint arXiv:1812.01718*, 2018.
- [63] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [64] Alex Krizhevsky and Geoffrey E. Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [65] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [66] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [67] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning*, pages 807–814, 2010.
- [68] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 448–456, 2015.
- [69] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8026–8037, 2019.
- [70] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations*, 2015.
- [71] Junnan Li, Richard Socher, and Steven C. H. Hoi. DivideMix: Learning with noisy labels as semi-supervised learning. In *Proceedings of the 8th International Conference on Learning Representations*, 2020.
- [72] Chongming Gao, Shijun Li, Wenqiang Lei, Jiawei Chen, Biao Li, Peng Jiang, Xiangnan He, Jiaxin Mao, and Tat-Seng Chua. KuaiRec: A fully-observed dataset and insights for evaluating recommender systems. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 540–550, 2022.
- [73] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*, pages 173–182, 2017.
- [74] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, pages 452–461, 2009.