

MULTI-CLASS CLASSIFICATION FROM SINGLE-CLASS DATA WITH CONFIDENCES

Anonymous authors

Paper under double-blind review

ABSTRACT

Can we learn a multi-class classifier from only *data of a single class*? We show that without any assumptions on the loss functions, models, and optimizers, we can successfully learn a multi-class classifier from only data of a single class with a rigorous consistency guarantee when *confidences* (i.e., the class-posterior probabilities for all the classes) are available. Specifically, we propose an empirical risk minimization framework that is loss-/model-/optimizer-independent. Instead of constructing a boundary between the given class and other classes, our method can conduct discriminative classification between all the classes even if no data from the other classes are provided. We further theoretically and experimentally show that our method can be Bayes-consistent with a simple modification even if the provided confidences are highly noisy. Then, we provide an extension of our method for the case where data from a subset of all the classes are available. Experimental results demonstrate the effectiveness of our methods.

1 INTRODUCTION

In supervised learning, the annotations of a huge number of instances may not be easily obtained in many practical applications due to the concerns including but not limited to time consumption, expenditure, and privacy preserving. For these reasons, many *weakly supervised learning* (WSL) frameworks (Zhou, 2018) have been studied in various scenarios recently, including *semi-supervised learning* (Chapelle et al., 2006; Zhu & Goldberg, 2009; Niu et al., 2013; Li & Zhou, 2015; Sakai et al., 2017; Li & Liang, 2019; Guo et al., 2020), *positive-unlabeled learning* (Elkan & Noto, 2008; du Plessis et al., 2014; 2015; Sansone et al., 2019), *unlabeled-unlabeled learning* (Lu et al., 2019; 2020), *noisy-label learning* (Natarajan et al., 2013; Han et al., 2018a;b; Zhang et al., 2019), *complementary and partial-label learning* (Ishida et al., 2017; Yu et al., 2018; Ishida et al., 2019; Feng et al., 2020a; Katsura & Uchida, 2020; Chou et al., 2020; Cour et al., 2011; Feng et al., 2020b; Lv et al., 2020), *similarity-based learning* (Bao et al., 2018; Shimada et al., 2019; Cao et al., 2021), and *positive-confidence learning* (Ishida et al., 2018).

In this paper, we investigate a novel weakly supervised learning scenario called ***Single-Class Confidence (SC-Conf) classification***, where only data of a single class annotated with the *confidences* (i.e., the class-posterior probabilities for all the classes) are required for training a *discriminative* multi-class classifier with *convergence guarantee on estimation error*, without any data from other classes. Such a WSL framework can be widespread in many real-world scenarios. For example, in the scenario of climatic disaster forecasting, in order to conduct ordinary supervised classification, we deploy sensors to collect data of normal climate and different kinds of meteorological disasters. However, under extreme climate conditions, e.g., typhoons, tornados, and snowstorms, the sensors may not work properly and we can only collect data from the distribution of normal climate. To construct a climatic disaster forecasting system that can discriminate different kinds of climates with only data of normal climate, we can ask meteorologists to annotate the data with a confidence score and train the system using our proposed framework with only data of normal climate. Another example is about market investigation. The investigators of each company aim to predict the consumers' tendency of purchasing the products of their companies and other competitors. However, due to privacy concerns, there may be data isolation, i.e., each company can only have access to their customers' information. To accomplish the investigation with the isolated data, the investigators can gather the purchase history of their customers and transform the purchase amount of each company




Figure 1: Illustration of the proposed method and related works. The points with dotted lines are not necessary in the training process and are only shown for illustration. We further show in Section 4 that even we have no information about the value of confidence scores, we can still conduct consistent SC-Conf classification only with the class with the maximum class-posterior probability: $y' = \operatorname{argmax}_{y \in \mathcal{Y}_s} p(y|\mathbf{x})$. We also show in Section 5 that it is possible to learn from a subset of all the classes.

into confidence between 0 and 1 by pre-processing. With the isolated data and confidence scores, the predictor can be trained by SC-Conf classification.

To apply the SC-Conf classification framework in more realistic scenarios, we show that with a small modification, the result of our method with extremely noisy confidence can be consistent with that with accurate confidence, i.e., our framework is noise-robust. To justify our claims, both infinite-sample and finite-sample analysis are provided.

Furthermore, we extend our method to the case where we can collect data from a subset of all the classes with confidences. For example, when conducting multi-class classification of climatic disasters including sandstorm, snowstorm, drought, and flood, we can train the classifier using data only with a coarse label ‘wind storm’ instead of collecting data from all four classes. Here we don’t have to specify the true labels of each instance, e.g., the further division of data with coarse label ‘wind storm’ into ‘sandstorm’ and ‘snowstorm’ is unnecessary. We refer to this framework as **Subset Confidence (Sub-Conf) classification** in the rest part of this paper.

Our contributions in this paper are three-fold:

- We provide an unbiased estimator of ordinary classification risk with only data of a single class and their confidences. An optimizer-independent empirical risk minimization (ERM) framework with no assumptions on loss functions and models is proposed. We further establish the estimation error bound to show the consistency of the proposed method.
- We show that if unlabeled samples are available, our proposed method is noise-robust and can be *classifier-consistent* and can converge to the optimal classifier with high probability even with extremely noisy confidence. Furthermore, we give a novel finite-sample convergence analysis on the misclassification rate of this method.
- An extension of our ERM framework to the case where data from a subset of all the classes with confidences are provided.

Extensive experiments with deep neural networks on both benchmark and real-world datasets are conducted for demonstrating the usefulness of our proposed methods.

2 RELATED WORK

In this section, we introduce related studies of the proposed SC-Conf classification framework. We briefly illustrate the proposed problem and related problems in Figure 1.

Multi-Class Classification For ordinary K -class multi-class classification problem, $\mathcal{X} \subset \mathbb{R}^d$ is the feature space and $\mathcal{Y} = [K]$ is the label space. Suppose each example $(\mathbf{x}, y) \in \mathcal{X} \times \mathcal{Y}$ is drawn independently from an unknown distribution with density $p(\mathbf{x}, y)$. To train a learning function for

multi-class classification $g : \mathcal{X} \rightarrow \mathbb{R}^K$, we have to minimize the classification risk below:

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

where $\ell(\cdot, \cdot) : \mathbb{R}^K \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is the multi-class loss function and $\mathbb{E}_{p(\mathbf{x}, y)}[\cdot]$ is the expectation on distribution with density $p(\mathbf{x}, y)$. The predicted label of \mathbf{x} is given as $f(\mathbf{x}) = \operatorname{argmax}_{y \in \mathcal{Y}} g_y(\mathbf{x})$, where $g_y(\mathbf{x})$ is the y -th element of $g(\mathbf{x})$. Since we are not aware of the joint density $p(\mathbf{x}, y)$, we collect identically and independently distributed examples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ and minimize the empirical risk (the sample mean of the losses of the collected examples) instead.

Multi-Positive and Unlabeled Learning. Multi-positive and unlabeled learning (Xu et al., 2017) is a WSL framework that can train multi-class classifiers using labeled data from $K - 1$ classes, unlabeled data collected from the distribution with density $p(\mathbf{x})$, and the class-prior probability of the unseen class. However, it is quite hard to cope with these requirements in real-world scenarios. On the other hand, our proposed SC-Conf classification framework only needs data of a single class and their confidence while it can still train multi-class classifiers.

One-Class Classification. One-class classification (Breunig et al., 2000; Tax & Duin, 1999; Schölkopf et al., 2001; Ruff et al., 2018) aims at ‘describing’ a given class with data of a single class rather than conducting discriminative classification. Furthermore, it cannot construct classification boundaries between all of the K classes with a finite-sample convergence guarantee since it has no access to the information of the rest $K-1$ classes. Compared with one-class classification, our proposed SC-Conf classification framework can construct discriminative multi-class classifiers with only data of a single class by utilizing the confidence scores. Due to the discussion above, it can be seen that the one-class classification methods are not suitable in the problem of multi-class classification.

Positive-Confidence Classification. In Ishida et al. (2018); Shinoda et al. (2020), Pconf classification is proposed to train discriminative binary classifiers only from positive examples and their positive-confidence (class-posterior probabilities of positive class). It can be regarded as a special case of our proposed SC-Conf classification when the number of classes is $K = 2$ since the class-posterior probability of the negative class can be immediately obtained when the Pconf is given. Nevertheless, it is still confined in the binary classification setting. When applied to multi-class classification, the Pconf framework can only discriminate the class regarded as the ‘positive’ class from the rest of all the classes, while the SC-Conf can conduct discriminative classification between all the classes, as shown in Figure 1.

3 SC-CONF CLASSIFICATION

In this section, we formulate the SC-Conf classification problem. In ordinary multi-class classification problem, we collect *i.i.d.* labeled data $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ from the joint distribution $p(\mathbf{x}, y)$ and conduct ERM (Vapnik, 1998) by minimizing the unbiased risk estimator of classification risk (1) constructed by the labeled data. However, under the SC-Conf classification framework, we are only given data from a **single** class y_s with confidence: $\mathcal{S} = \{\mathbf{x}_i, \mathbf{r}_i\}_{i=1}^n$, where \mathbf{x}_i is an instance drawn independently from the distribution of a single class with density $p(\mathbf{x}|y = y_s)$ and \mathbf{r}_i is a confidence vector given by $\{p(y|\mathbf{x}_i)\}_{y=1}^K$ ¹. Based on these information, we show how to construct an unbiased risk estimator and conduct ERM without any instance from the other $K - 1$ classes in this section.

3.1 UNBIASED RISK ESTIMATOR

We denote the class-prior probability of class y_s as π_{y_s} . Let $\mathbf{r}(\mathbf{x}) = \{p(y|\mathbf{x})\}_{y=1}^K$ and $r^i(\mathbf{x}) = p(y = i|\mathbf{x})$. $\mathbb{E}_{p(\mathbf{x}|y_s)}[\cdot]$ is the expectation on the distribution with density $p(\mathbf{x}|y_s)$. The following theorem shows that the classification risk can be recovered with the supervision above:

Theorem 1. *The multi-class classification risk (1) can be recovered from data drawn from a single class if their confidence scores are given:*

$$R(g) = \pi_{y_s} \mathbb{E}_{p(\mathbf{x}|y_s)} \left[\sum_{y=1}^K \frac{r^y(\mathbf{x})}{r^{y_s}(\mathbf{x})} \ell(g(\mathbf{x}), y) \right], \quad (2)$$

where $p(y_s|\mathbf{x}) > 0$ for all \mathbf{x} in the support of the distribution with density $p(\mathbf{x})$.

¹Note that in real-world applications, the equation $\mathbf{r}_i = p(y|\mathbf{x}_i)$ may not hold due to the noisy supervision. We analyze the effect of noisy confidence in Section 4.

The proof is shown in Appendix A. A sketch of the proof is that (2) is essentially duplicating (\mathbf{x}, \mathbf{r}) into $(\mathbf{x}, p(y = 1|\mathbf{x}))$, $(\mathbf{x}, p(y = 2|\mathbf{x}))$, ..., $(\mathbf{x}, p(y = K|\mathbf{x}))$ and then applying importance weighting to adjust the difference between the distribution of a single class and all the classes. It is noteworthy that the requirement that $p(y_s|\mathbf{x}) > 0$ is widespread and used in the previous works (Ishida et al., 2018; Shinoda et al., 2020). It can also be easily satisfied in many practical applications. For example, when the market investigators of a leading company analyzing the tendency of their potential consumers, they will find that almost all of the consumers (i.e., \mathbf{x}) have interest of their products (i.e., y_s) since their company plays a dominant role in the market. According to Theorem 1, we can see that Eq. (2) is an equivalent risk expression of Eq. (1). Then we can directly get the unbiased risk estimator of Eq. (2):

$$\hat{R}_{SC}(\mathbf{g}) = \frac{\pi_{y_s}}{n} \sum_{i=1}^n \sum_{y=1}^K \frac{r_i^y}{r_i^{y_s}} \ell(\mathbf{g}(\mathbf{x}_i), y). \quad (3)$$

Since Eq. (2) does not include the expectation over the data of the other $K - 1$ classes, only data from the single class y_s are used in Eq. (3) for constructing the unbiased risk estimator. Then the following work is to conduct ERM by minimizing the unbiased risk estimator in Eq. (3). Since the class-prior probability π_{y_s} can be safely ignored without changing the result of ERM, we do not have to estimate π_s in the process of SC-Conf classification.

3.2 ESTIMATION ERROR BOUND

Here we analyze the consistency of the proposed unbiased risk estimator $\hat{R}_{SC}(\mathbf{g})$ in (3) and give the convergence rate of its estimation error. First of all, let $\mathcal{G} = [\mathcal{G}_y]_{y=1}^K$ be a class of K -valued functions that we consider in ERM. Assume that there is $C_g > 0$ that $\sup_{\mathbf{g} \in \mathcal{G}} \|\mathbf{g}\|_\infty \leq C_g$. The loss function $\ell(\mathbf{g}(\mathbf{x}), y)$ is Lipschitz continuous for all $\|\mathbf{g}\|_\infty \leq C_g$ with Lipschitz constant $L_\ell > 0$ and upper-bounded by $C_\ell > 0$. At the end, we assume that the confidence of the given single class is not too small that there exists $C_r > 0$ such that $p(y_s|\mathbf{x}) > C_r$ holds almost surely. Though we can omit the $p(y_s|\mathbf{x})$ and analyze the biased version of $\hat{R}_{SC}(\mathbf{g})$, for simplicity, only the unbiased version is analyzed here. Denote by $\hat{\mathbf{g}}_{SC}$ the minimizer of $\hat{R}_{SC}(\mathbf{g})$ in (3) and \mathbf{g}^* the minimizer of (1). Then the estimation error is given by $R(\hat{\mathbf{g}}) - R(\mathbf{g}^*)$. The following theorem gives the estimation error bound:

Theorem 2. For any $0 < \delta < 1$, with probability at least $1 - \delta$:

$$R(\hat{\mathbf{g}}_{SC}) - R(\mathbf{g}^*) \leq \frac{4\sqrt{2}\pi_{y_s}L_\ell}{C_r} \sum_{y=1}^K \mathfrak{R}_n(\mathcal{G}_y) + 2\pi_{y_s}C_\ell \sqrt{\frac{\log \frac{2}{\delta}}{2n}}, \quad (4)$$

where $\mathfrak{R}_n(\mathcal{G}_y)$ is the Rademacher complexity (Bartlett & Mendelson, 2001) of \mathcal{G}_y on i.i.d. samples with size n drawn from a single class $p(\mathbf{x}|y_s)$.

The definition of Rademacher complexity and the proof is given in the Appendix B. Generally, the Rademacher complexity $\mathfrak{R}_n(\mathcal{G}_y)$ has an upper bound of $C_{\mathcal{G}_y}/\sqrt{n}$ (Golowich et al., 2018), where $C_{\mathcal{G}_y}$ is a positive number determined by the function class \mathcal{G}_y . We can see that the risk of empirically optimal classifier $R(\hat{\mathbf{g}}_{SC})$ converges to that of the optimal classifier $R(\mathbf{g}^*)$ in the rate of $O_p(1/\sqrt{n})$, which is the optimal parametric rate in probability and cannot be improved without additional assumptions (Mendelson, 2008).

4 CONSISTENT NOISE-ROBUST METHOD FOR SC-CONF CLASSIFICATION

In Section 3, it can be seen that the confidences play an important role in the construction of the unbiased risk estimator (3). However, in practical applications, the confidences $\{r_i\}_{i=1}^n$ are often inaccurate, and we can only get the noisy confidence scores $\{\tilde{r}_i\}_{i=1}^n$ generated by a corrupted distribution. Since $\tilde{r}_i \neq r_i$, the risk estimator proposed in Section 3 is biased in general if we simply plug the noisy confidence into (3). As a result, the learning guarantee in the previous section no longer holds and thus the use of ERM is not reasonable. Can we obtain a provably consistent multi-class classifier with data of a single class and extremely noisy confidence scores? Surprisingly, in this section, we give a positive answer to this question.

We propose a refined method called **Noise-Robust Single-Class Confidence** (NoRSC-Conf) classification, where only data of a single class, unlabeled data, and noisy confidence scores are provided. Then we show that this method is *classifier-consistent* (Yu et al., 2018; Feng et al., 2020b; Lv et al.,

2020): the predictions of the classifier generated from noisy supervision is still infinite-sample consistent to those of the optimal classifier of classification risk $R(\mathbf{g})$. Based on this result, we establish the ERM framework for NoRSC-Conf classification. Since there is a gap between the convergence on the corrupted distribution and the original one, we cannot directly analyze the finite-sample consistency of NoRSC-Conf classification with traditional techniques on bounding the uniform convergence. To justify the use of ERM, we exploit the connection between the noisy and original distribution and give a novel finite-sample analysis on the misclassification rate of NoRSC-Conf classification. Finally, we show that the weight $\phi(\mathbf{x})$ used in the construction of NoRSC-Conf classification can be estimated from unlabeled data and data of a single class by applying the density ratio matching-based method (Sugiyama et al., 2012) and give its estimation error bound.

4.1 FORMULATION OF NORSC-CONF CLASSIFICATION

In NoRSC-Conf classification, we have *i.i.d.* unlabeled data $\mathcal{S}_u = \{\mathbf{x}_i^u\}_{i=1}^{n_u}$ drawn from the distribution with marginal density $p(\mathbf{x})$, data of a single class with noisy confidence scores $\tilde{\mathcal{S}} = \{\mathbf{x}_i, \tilde{r}_i\}_{i=1}^n$. Here the generation of data of a single class $\{\mathbf{x}_i\}_{i=1}^n$ is the same with that in the previous section, and the noisy confidence $\{\tilde{r}_i\}_{i=1}^n$ are generated by a corrupted distribution with density $\tilde{p}(y|\mathbf{x})$.

Suppose the optimal classifier $\tilde{\mathbf{g}}^{\text{bayes}} = \operatorname{argmin}_{\mathbf{g} \text{ measurable}} \tilde{R}(\mathbf{g})$ is in the function class \mathcal{G} , which is a widely used assumption in classifier-consistency analysis, e.g., the Assumption 1 in Yu et al. (2018). Straightforwardly, we give the the formulation and show the infinite-sample consistency of the NoRSC-Conf classification method:

Theorem 3. *For each instance \mathbf{x} , if $\operatorname{argmax}_{y \in \mathcal{Y}} \tilde{p}(y|\mathbf{x}) = \operatorname{argmax}_{y \in \mathcal{Y}} p(y|\mathbf{x})$ and the classification-calibrated losses (e.g., softmax cross entropy loss, mean squared error) (Tewari & Bartlett, 2007) are used, the following reweighed risk formulation is classifier-consistent:*

$$\tilde{R}(\mathbf{g}) = \mathbb{E}_{p(\mathbf{x}|y_s)} [\phi(\mathbf{x}) \sum_{y=1}^K \tilde{r}^y(\mathbf{x}) \ell(\mathbf{g}(\mathbf{x}), y)], \quad (5)$$

where $\phi(\mathbf{x}) = \frac{p(\mathbf{x})}{p(\mathbf{x}|y_s)}$, $\tilde{r}^y(\mathbf{x}) = \tilde{p}(y|\mathbf{x})$, and $p(y_s|\mathbf{x}) > 0$ for all \mathbf{x} in the support of the distribution with density $p(\mathbf{x})$.

This assumption is realistic in various scenarios, e.g., when the generation process of supervision is corrupted by some class-conditional noise (Zhang et al., 2021) or some noise that meets the row-diagonally dominant conditional (Gui et al., 2021). An intuitive explanation is that the decision boundary of the Bayes-optimal classifier is only determined by the class with the maximum class-posterior probability of each instance. Then we only have to adjust the distribution difference between the distribution of a single class and all the classes. The detailed proof is shown in Appendix C. With regard to the assumption on the noisy confidence, Theorem 3 only requires that for each instance \mathbf{x} , the class with the largest posterior probability in the original distribution should still have the largest value among all the classes in the noisy distribution. Notice that since there is no assumption on the exact value of the confidence scores, the confidence scores can be of arbitrary values in the K -dimensional probability simplex as long as the assumption above is satisfied. This assumption is realistic in practice. For example, when asking the annotators for labeling the data of a single class with confidence scores, they may not be sure about the exact value of the true class-posterior probability of each class. However, it can be easier for them to identify the class of $\operatorname{argmax}_{y \in \mathcal{Y}} p(y|\mathbf{x})$ and give the highest but not necessarily the accurate confidence score to this class. Though this kind of confidence score is not accurate in general, according to Theorem 3, it is meaningful in the setting of NoRSC-Conf classification and can be used for generating consistent classifiers.

Denote by $\tilde{\mathbf{g}}^*$ the minimizer of (5), then Theorem 3 means that $\tilde{\mathbf{g}}^*$ and \mathbf{g}^* have the same prediction, i.e., $\operatorname{argmax}_{y \in \mathcal{Y}} \tilde{g}_y^*(\mathbf{x}) = \operatorname{argmax}_{y \in \mathcal{Y}} g_y^*(\mathbf{x})$ for all the \mathbf{x} . Thus the prediction of the minimizer of (5) is infinite-sample consistent to that of the ordinary classification risk (1). Based on this result and $\phi_i = \phi(\mathbf{x}_i)$, it is common practice to conduct ERM by minimizing the unbiased estimator of (5):

$$\hat{R}(\mathbf{g}) = \frac{1}{n} \sum_{i=1}^n \phi_i \sum_{y=1}^K \tilde{r}_i^y \ell(\mathbf{g}(\mathbf{x}_i), y). \quad (6)$$

Denote by $\hat{\mathbf{g}}$ the minimizer of the empirical risk (6). Using the same technique of bounding the uniform convergence as in Section 3.2, we can get the finite-sample analysis of the result: $\hat{R}(\hat{\mathbf{g}}) \rightarrow \hat{R}(\tilde{\mathbf{g}}^*)$. Nevertheless, in the scenario of multi-class classification, we are concerning about the problem

that if the misclassification risk of $\hat{\mathbf{g}}$ on the original distribution with density $p(\mathbf{x}, y)$ converges to the optimal one, i.e., if $R_{01}(\hat{\mathbf{g}}) \rightarrow R_{01}(\mathbf{g}_{01}^*)^2$, where $R_{01}(\mathbf{g}) = \mathbb{E}_{p(\mathbf{x}, y)} [1(\arg\max_{y' \in \mathcal{Y}} g_{y'}(\mathbf{x}) \neq y)]$ is the equivalent expression of misclassification rate of classifier \mathbf{g} , and $\mathbf{g}_{01}^* = \arg\min_{\mathbf{g} \in \mathcal{G}} R_{01}(\mathbf{g})$ is the optimal classifier of misclassification rate. Since there is difference between the noisy and original distribution, we cannot get the convergence result of $R_{01}(\hat{\mathbf{g}})$ directly from the fact that $\tilde{R}(\hat{\mathbf{g}}) \rightarrow \tilde{R}(\tilde{\mathbf{g}}^*)$. Thus the finite-sample analysis of the misclassification rate of $\hat{\mathbf{g}}$ is non-trivial.

In the following section, we will give the answer to this question by exploiting the connection between $\tilde{p}(y|\mathbf{x})$ and $p(y|\mathbf{x})$ and give a finite-sample analysis on the misclassification rate of $\hat{\mathbf{g}}$.

4.2 FINITE-SAMPLE ANALYSIS ON THE MISCLASSIFICATION RATE

In this section, we give a non-trivial finite-sample analysis on the misclassification rate of $\hat{\mathbf{g}}$: $R_{01}(\hat{\mathbf{g}}) - R_{01}(\mathbf{g}_{01}^*)$. To begin with, we first give the convergence analysis on $\tilde{R}(\hat{\mathbf{g}}) \rightarrow \tilde{R}(\tilde{\mathbf{g}}^*)$.

Lemma 1. *Based on the assumptions in previous parts, for any $0 < \delta < 1$, with probability at least $1 - \delta$:*

$$\tilde{R}(\hat{\mathbf{g}}) - \tilde{R}(\tilde{\mathbf{g}}^*) \leq \frac{4\sqrt{2}L_\ell}{C_r} \sum_{y=1}^K \mathfrak{R}_n(\mathcal{G}_y) + \frac{2C_\ell}{C_r} \sqrt{\frac{\log \frac{2}{\delta}}{2n}}, \quad (7)$$

where $\mathfrak{R}_n(\mathcal{G}_y)$ is the Rademacher complexity of \mathcal{G}_y on i.i.d. samples with size n drawn from a single class $p(\mathbf{x}|y_s)$.

Combining this lemma with the Corollary 26 in (Zhang, 2004), we can directly bound the following term w.r.t. 0-1 loss: $\tilde{R}_{01}(\hat{\mathbf{g}}) - \tilde{R}_{01}(\tilde{\mathbf{g}}^*)$:

Lemma 2. *If the loss function is classification-calibrated, then there exists a concave function ξ on $[0, +\infty)$ such that $\xi(0) = 0$ and $\xi(\delta) \rightarrow 0$ as $\delta \rightarrow 0^+$. Furthermore, for any $0 < \delta < 1$, with probability at least $1 - \delta$:*

$$\tilde{R}_{01}(\hat{\mathbf{g}}) - \tilde{R}_{01}(\tilde{\mathbf{g}}^*) \leq \xi(\tilde{R}(\hat{\mathbf{g}}) - \tilde{R}(\tilde{\mathbf{g}}^*)). \quad (8)$$

where $\tilde{R}_{01}(\mathbf{g})$ is the misclassification rate of \mathbf{g} on the noisy distribution with density $\tilde{p}(\mathbf{x}, y) = \tilde{p}(y|\mathbf{x})p(\mathbf{x})$. The validity of the density $\tilde{p}(\mathbf{x}, y)$ is shown in the proof of Theorem 3.

The lemma above shows that $\tilde{R}_{01}(\hat{\mathbf{g}}) \rightarrow \tilde{R}_{01}(\tilde{\mathbf{g}}^*)$ in probability as $n \rightarrow +\infty$. However, this conclusion is still not informative enough: since there is a gap between the noisy density $\tilde{p}(\mathbf{x}, y)$ and the original distribution $p(\mathbf{x}, y)$, the conclusion $\tilde{R}_{01}(\hat{\mathbf{g}}) \rightarrow \tilde{R}_{01}(\tilde{\mathbf{g}}^*)$ in probability cannot directly give an answer to the question that if the risk of the empirical minimizer of NoRSC-Conf classification $R_{01}(\hat{\mathbf{g}})$ can converge to the optimal misclassification rate $R_{01}(\mathbf{g}_{01}^*)$ on the original distribution in probability. To answer this question, we have to make use of the connection between $\tilde{p}(\mathbf{x}, y)$ and $p(\mathbf{x}, y)$. For each $\mathbf{x} \in \mathcal{X}$, let $\Delta(\mathbf{x}) = \tilde{p}(y|\mathbf{x}) - \tilde{p}(y'|\mathbf{x})$, where y and y' are the classes with the largest and the second largest posterior possibilities, respectively. Then we have the following conclusion:

Theorem 4. *Suppose $\inf_{\mathbf{x} \in \mathcal{X}} \Delta(\mathbf{x}) > 0$, then for any $0 < \delta < 1$, with probability at least $1 - \delta$:*

$$R_{01}(\hat{\mathbf{g}}) - R_{01}(\mathbf{g}_{01}^*) \leq \frac{1}{\inf_{\mathbf{x} \in \mathcal{X}} \Delta(\mathbf{x})} \xi(\tilde{R}(\hat{\mathbf{g}}) - \tilde{R}(\tilde{\mathbf{g}}^*)). \quad (9)$$

We prove this conclusion in Appendix D. Combining Lemma 1, Theorem 4, and the property of function ξ , we can finally say that with the increasing of sample size n , the misclassification rate of $\hat{\mathbf{g}}$ converges to the optimal misclassification rate $R_{01}(\mathbf{g}_{01}^*)$ in probability. According to this conclusion, we can see that the minimizer of the empirical risk 6 possesses finite-sample consistency, which justifies the use of ERM.

Remark 1. A representative scenario is that only the class with maximum class-posterior probability is identified for each data point. Then the noisy confidences are given in the form of one-hot codes: $\tilde{\mathbf{r}}(\mathbf{x}) = [0, \dots, 1, \dots, 0]$, where the only non-zero element 1 is on the y_{sc} -th position and y_{sc} is the class with maximum class-posterior probability. In this scenario, almost all the details of the confidence of each class are lost and the given confidences are extremely noisy. It can be seen that in this scenario, $\inf_{\mathbf{x} \in \mathcal{X}} \Delta(\mathbf{x}) = 1$. According to Theorem 4, the result of our NoRSC-Conf classification can still converge to the Bayes-optimal classifier in probability.

²According to the property of classification-calibrated loss, the minimizer \mathbf{g}^* of $R(\mathbf{g})$ with surrogate loss ℓ is also the minimizer of misclassification rate $R_{01}(\mathbf{g})$ among all the measurable functions, where $1(\cdot)$ is the indicator function.

4.3 EFFICIENT ESTIMATION OF DENSITY RATIO $\phi(\mathbf{x})$ WITH UNLABELED DATA

Though the infinite-sample consistency and the finite-sample consistency of our NoRSC-Conf classification have been established, there is still an important problem to be solved: the estimation of the weight term $\phi(\mathbf{x})$. In this section, we show that additional information of the class labels is not required in the estimation of the weight term $\phi(\mathbf{x})$ and only the data from only a single class and unlabeled data are needed.

Let $\phi(\mathbf{x})$ be the true density ratio $p(\mathbf{x})/p(\mathbf{x}|y_s)$ and $\hat{\phi}(\mathbf{x})$ be the estimated one. We can empirically estimate $\phi(\mathbf{x})$ by the density ratio matching method (Sugiyama et al., 2012). Here, we give the definition of *Bregman divergence*, which measures the discrepancy between two density ratio models:

Definition 1. (Sugiyama et al., 2012) For any differentiable and strictly convex function $\eta: \mathbb{R} \rightarrow \mathbb{R}$, $\nabla\eta(t)$ is the subgradient of η . The Bregman divergence of η between the true density ratio $\phi(\mathbf{x})$ and the estimated density ratio $\hat{\phi}(\mathbf{x})$ is given as:

$$B_\eta(\phi\|\hat{\phi}) = \int p(\mathbf{x}|y_s)\nabla\eta(\hat{\phi}(\mathbf{x}))\hat{\phi}(\mathbf{x})d\mathbf{x} - \int p(\mathbf{x}|y_s)\eta(\hat{\phi}(\mathbf{x}))d\mathbf{x} - \int p(\mathbf{x})\nabla\eta(\hat{\phi}(\mathbf{x}))d\mathbf{x},$$

and its unbiased estimator is given as:

$$\hat{B}_\eta(\phi\|\hat{\phi}) = \frac{1}{n} \sum_{i=1}^n \nabla\eta(\hat{\phi}(\mathbf{x}_i))\hat{\phi}(\mathbf{x}_i) - \frac{1}{n} \sum_{i=1}^n \eta(\hat{\phi}(\mathbf{x}_i)) - \frac{1}{n_u} \sum_{i=1}^{n_u} \nabla\eta(\hat{\phi}(\mathbf{x}_i^u)). \quad (10)$$

Since we are given data from the single class y_s and the unlabeled data, we can estimate the true density ratio by minimizing the empirical Bregman divergence (10). Denote by Φ the function class of ratio model and $\hat{\phi}^* = \operatorname{argmin}_{\hat{\phi} \in \Phi} \hat{B}_\eta(\phi\|\hat{\phi})$ and assume that the true ratio model $\phi \in \Phi$, we have the following estimation error bound:

Theorem 5. Based on the assumptions above, for any $0 < \delta < 1$, with probability at least $1 - \delta$:

$$B_\eta(\phi\|\hat{\phi}^*) \leq C_1 \mathfrak{R}_n(\Phi) + C_2 \mathfrak{R}_{n_u}^u(\Phi) + M_1 \sqrt{\frac{\ln \frac{4}{\delta}}{2n}} + M_2 \sqrt{\frac{\ln \frac{4}{\delta}}{2n_u}}$$

where $\mathfrak{R}_n(\Phi)$ is the Rademacher complexity of function class Φ on i.i.d. samples with size n drawn from a single class $p(\mathbf{x}|y_s)$, $\mathfrak{R}_{n_u}^u(\Phi)$ is the Rademacher complexity of function class Φ on i.i.d. samples with size n_u drawn from marginal distribution with density $p(\mathbf{x})$. Meanwhile, C_1, C_2, M_1, M_2 are constants.

The proof is given in Appendix E. This estimation error bound guarantees that the estimated ratio $\hat{\phi}$ converges to the true one in the rate of $\mathcal{O}_p(1/\sqrt{n_u} + 1/\sqrt{n})$. In other words, our estimated ratio model will be more accurate as the numbers of data from the single class y_s and unlabeled data increases.

5 SUB-CONF CLASSIFICATION

In many practical situations, the data may not be collected from only a single class $y_s \in \mathcal{Y}$ but from a subset $\mathcal{Y}_s \subset \mathcal{Y}$ of all the classes. Here we give the data generation process of this kind of data. Suppose we have data drawn from a subset of all the classes and their confidence: $\{\mathbf{x}_i, r_i\}_{i=1}^n$, where the confidence is the same with that in Section 3 and the data $\{\mathbf{x}_i\}_{i=1}^n$ are i.i.d. samples drawn from an unknown distribution with conditional density $p(\mathbf{x}|y \in \mathcal{Y}_s)$. Then we can get an equivalent risk expression of the classification risk (1), which is similar to the formulation (2):

Theorem 6. The multi-class classification risk (1) can be recovered from data drawn from a subset of the collection of all the classes, i.e., $\mathcal{Y}_s \subset \mathcal{Y}$, if their confidence scores are given:

$$R(\mathbf{g}) = \pi_{\mathcal{Y}_s} \mathbb{E}_{p(\mathbf{x}|y \in \mathcal{Y}_s)} \left[\sum_{y=1}^K \frac{r^y(\mathbf{x})}{r^{\mathcal{Y}_s}(\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) \right], \quad (11)$$

where $\pi_{\mathcal{Y}_s} = \sum_{y_s \in \mathcal{Y}_s} \pi_{y_s}$ and $r^{\mathcal{Y}_s}(\mathbf{x}) = \sum_{y_s \in \mathcal{Y}_s} r^{y_s}(\mathbf{x})$.

The proof is shown in Appendix F. With this conclusion, we can approximate the classification risk (1) with data from a subset of all the classes as:

$$\hat{R}_{\text{sub}}(\mathbf{g}) = \frac{\pi_{\mathcal{Y}_s}}{n} \sum_{i=1}^n \sum_{y=1}^K \frac{r_i^y}{r_i^{\mathcal{Y}_s}} \ell(\mathbf{g}(\mathbf{x}_i), y), \quad (12)$$

where $r_i^{\mathcal{Y}_s} = \sum_{y_s \in \mathcal{Y}_s} r_i^{y_s}$. Then it is routine to conduct ERM based on the unbiased risk estimator (12). When $|\mathcal{Y}_s| = 1$, the Sub-Conf classification is equivalent to the SC-Conf classification, which indicates that Sub-Conf classification is the generalization of SC-Conf classification. Similarly, the class-prior probability $\pi_{\mathcal{Y}_s}$ can be eliminated in the process of ERM.

Remark 2. A potential way of constructing equivalent risk expression with data from classes $\mathcal{Y}_s \subset \mathcal{Y}$ is the convex combination of SC-Conf classification:

$$R(\mathbf{g}) = \sum_{y_s \in \mathcal{Y}_s} \alpha_{y_s} \pi_{y_s} \mathbb{E}_{P(\mathbf{x}|y_s)} \left[\sum_{y=1}^K \frac{r^y(\mathbf{x})}{r^{y_s}(\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) \right], \text{ s.t. } \sum_{y_s \in \mathcal{Y}_s} \alpha_{y_s} = 1, \alpha_{y_s} \geq 0,$$

which is a straightforward extension of SC-Conf classification. However, this method requires that we have the exact label $y_s \in \mathcal{Y}_s$ for each instance. In contrast, it can be seen from the formulation of the empirical risk (12) that the exact class label for each instance is not necessary in our Sub-Conf classification method, which shows that our Sub-Conf classification method is non-trivial.

It can be seen that Sub-Conf classification and SC-Conf classification share the nearly identical risk formulation and data generation process and we can get a similar conclusion for Sub-Conf classification on its convergence analysis by substituting y_s with \mathcal{Y}_s . Given space limitations, we omit the convergence analysis of Sub-Conf classification and show the specific formulation of Noise-Robust Sub-Conf classification in the Appendix G.

6 EXPERIMENTS

In this section, we experimentally show the usefulness of our proposed methods for training deep models on two benchmark datasets. The implementation is conducted on NVIDIA GeForce RTX 3090 GPUs based on Pytorch (Paszke et al., 2019) and Sklearn (Pedregosa et al., 2011).

Baselines: In the experiments, SC-Conf, NoRSC-Conf, and Sub-Conf are short for ERM with risk estimators (3), (6), and (12). We compare our methods to a simple weighted classification estimator:

$$\hat{R}_w(\mathbf{g}) = \frac{1}{n} \sum_{i=1}^n \sum_{y=1}^K r_i^y \ell(\mathbf{g}(\mathbf{x}_i), y). \quad (13)$$

This binary version of this estimator was used as a baseline in Pconf classification (Tshida et al., 2018). We regard ERM using this risk estimator with Weighted in the following parts. Though this estimator is a natural idea of utilizing the confidences when the data is collected from all the classes, it is biased generally in our setting where only data of a single class or a subset of all the classes are available. We also offer the result of learning with fully-supervised data. The details of the used datasets and our experimental setups on them are shown in this section.

Datasets: We evaluate the performance of the proposed methods and baselines on *Fasion-MNIST* dataset (Xiao et al., 2017) and *CIFAR-10* dataset (Krizhevsky et al., 2009). In the experiments, we use data of a single class or a subset of all the classes to train a classifier that can conduct classification successfully on all the classes. The detailed information of the used datasets are listed in the supplementary materials.

Setup: We briefly introduce the setting of our experiments and the detailed statistics are shown in the supplementary materials. We train the proposed methods and baseline methods with 3-layer multi-layer perceptron and DenseNet-161 (Huang et al., 2017) with softmax cross-entropy loss on Fashion-MNIST and CIFAR-10, respectively. Adam (Kingma & Ba, 2015) is used as the optimization algorithm. The validation accuracy of the used methods are calculated according to their empirical risk estimators on a validation set consisted of SC/Sub-Conf data. Then we do not have to collect additional labeled data for validation. Though we can ask annotators for the values of confidences in real-world applications, we simulate the confidences here by a probabilistic model. For generating the confidences, we use the separated labeled dataset for training a probabilistic model with the same loss function as in the last paragraph and this dataset is separated from any other process of experiments. Since the minimizer of cross-entropy loss is a good estimator of the class-posterior probability (Yu et al., 2018; Lv et al., 2020; Feng et al., 2020b), we use the output of the training and validation sets after a softmax layer as their confidences. As to the noisy confidences, we consider the scenario in Remark 1. Since most of the elements are zero in this scenario, the SC/Sub-Conf methods are not applicable here since the denominators in (3) and (12) may be zeros. We conduct experiments on the proposed NoRSC-Conf method and the baseline weighted method to show the robustness of NoRSC-Conf against extreme noise.

Results: The experimental results are shown in Table 1 and 2. As we can see, when accurate confidences are used, SC/Sub-Conf outperforms NoRSC-Conf and Weighted in almost all the cases. When the size of the subset of all the classes increases, the performance of SC/Sub-Conf is even

Table 1: Mean and standard deviation of the classification accuracy over 10 trials for the Fashion-MNIST dataset. The proposed methods were compared with the baseline Weighted method and fully-supervised method, with different classes used for training. The best and equivalent methods are shown in bold based on the 5% t-test, excluding fully-supervised method.

Used Classes		SC/Sub-Conf	NoRSC-Conf	Weighted	Supervised
Coat	Accurate	54.66±2.17	49.72±2.35	49.63±1.82	80.13±2.75
	Noisy	-±-	49.64±3.06	49.29±3.87	
Sandal	Accurate	56.02±4.55	45.65±6.42	44.59±7.98	
	Noisy	-±-	44.50±5.19	40.28±4.85	
Shirt	Accurate	71.44±4.77	60.26±7.45	57.90±7.46	
	Noisy	-±-	59.16±4.51	54.23±4.33	
Bag	Accurate	71.29±2.99	66.04±2.03	68.31±1.56	
	Noisy	-±-	63.19±4.24	60.67±2.70	
Trouser & Pullover	Accurate	62.05±2.92	54.25±4.17	53.86±4.68	
	Noisy	-±-	53.74±2.73	50.56±3.87	
Dress & Sneaker & Ankle Boot	Accurate	76.96±2.03	74.64±2.38	73.06±2.43	
	Noisy	-±-	72.14±2.57	71.42±3.23	

Table 2: Mean and standard deviation of the classification accuracy over 5 trials for the CIFAR-10 dataset. The proposed methods were compared with the baseline Weighted method and fully-supervised method, with different classes used for training. The best and equivalent methods are shown in bold based on the 5% t-test, excluding fully-supervised method.

Used Classes		SC/Sub-Conf	NoRSC-Conf	Weighted	Supervised
Dog	Accurate	53.10±0.93	51.52±0.76	51.60±0.98	65.62±0.84
	Noisy	-±-	41.77±0.70	39.94±0.74	
Airplane & Dog	Accurate	56.73±0.42	55.69±0.68	54.47±1.21	
	Noisy	-±-	46.47±0.59	45.42±0.83	
Airplane & Deer & Ship	Accurate	57.12±1.63	53.69±1.70	54.68±0.64	
	Noisy	-±-	51.11±0.65	48.96±1.00	
Airplane & Cat & Frog & Truck	Accurate	60.11±0.27	58.95±0.61	57.47±1.26	
	Noisy	-±-	54.06±0.47	52.56±0.85	
Airplane & Cat & Frog & Ship & Truck	Accurate	63.25±1.53	61.19±0.95	59.89±1.43	
	Noisy	-±-	58.95±1.05	54.42±2.38	

comparable to that of fully-supervised learning. Notice that since we only have data from a single class or a subset of all the classes, the training samples used for SC/Sub-Conf learning are far less than the whole training set. For a fair comparison, the number of training samples used for fully-supervised learning is reduced to that of the subset with the most samples. When only the noisy confidences are given, NoRSC-Conf outperforms the weighted classification baseline in eight cases and is comparable to it in three cases. This result shows that NoRSC-Conf can alleviate the effect of noisy supervision and identify the Bayes-optimal classifier, which aligns with our theoretical analysis in Section 4.

7 CONCLUSION

In this paper, we propose a novel weakly supervised learning setting and effective algorithms for provably consistent multi-class classification from data of a single class or a subset of all the classes equipped with confidences. We make three key contributions in this paper. Firstly, we propose an unbiased risk estimator for multi-class classification from data of a single class with confidences and provided the estimation error analysis. Secondly, we theoretically show that the proposed method can be robust to extreme noise and converge to the Bayes-optimal classifier in probability, and both infinite and finite-sample analyses on the misclassification rate are given. Finally, we extend our method to the case where data from a subset of all the classes are available. The experimental results demonstrate the usefulness of our algorithms.

REFERENCES

- Han Bao, Gang Niu, and Masashi Sugiyama. Classification from pairwise similarity and unlabeled data. In *ICML*, pp. 461–470, 2018.
- Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. In *COLT*, pp. 224–240, 2001.
- Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. LOF: Identifying density-based local outliers. In *SIGMOD*, pp. 93–104, 2000.
- Yuzhou Cao, Lei Feng, Yitian Xu, Bo An, Gang Niu, and Masashi Sugiyama. Learning from similarity-confidence data. *CoRR*, abs/2102.06879, 2021.
- Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien (eds.). *Semi-Supervised Learning*. The MIT Press, 2006.
- Yu-Ting Chou, Gang Niu, Hsuan-Tien Lin, and Masashi Sugiyama. Unbiased risk estimators can mislead: A case study of learning with complementary labels. In *ICML*, pp. 1929–1938, 2020.
- Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. *CoRR*, abs/1812.01718, 2018.
- Timothée Cour, Benjamin Sapp, and Ben Taskar. Learning from partial labels. *J. Mach. Learn. Res.*, 12:1501–1536, 2011.
- Marthinus Christoffel du Plessis, Gang Niu, and Masashi Sugiyama. Analysis of learning from positive and unlabeled data. In *NeurIPS*, pp. 703–711, 2014.
- Marthinus Christoffel du Plessis, Gang Niu, and Masashi Sugiyama. Convex formulation for learning from positive and unlabeled data. In *ICML*, pp. 1386–1394, 2015.
- Charles Elkan and Keith Noto. Learning classifiers from only positive and unlabeled data. In *KDD*, pp. 213–220, 2008.
- Lei Feng, Takuo Kaneko, Bo Han, Gang Niu, Bo An, and Masashi Sugiyama. Learning from multiple complementary labels. In *ICML*, pp. 3072–3081, 2020a.
- Lei Feng, Jiaqi Lv, Bo Han, Miao Xu, Gang Niu, Xin Geng, Bo An, and Masashi Sugiyama. Provably consistent partial-label learning. In *ICML*, 2020b.
- Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In *COLT*, pp. 297–299, 2018.
- Xian-Jin Gui, Wei Wang, and Zhang-Hao Tian. Towards understanding deep learning from noisy labels with small-loss criterion. In *IJCAI*, pp. 2469–2475, 2021.
- Lan-Zhe Guo, Zhi Zhou, and Yu-Feng Li. Record: Resource constrained semi-supervised learning under distribution shift. In *KDD*, pp. 1636–1644, 2020.
- Bo Han, Jiangchao Yao, Gang Niu, Mingyuan Zhou, Ivor W. Tsang, Ya Zhang, and Masashi Sugiyama. Masking: A new perspective of noisy supervision. In *NeurIPS*, pp. 5841–5851, 2018a.
- Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor W. Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, pp. 8536–8546, 2018b.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pp. 463–469, 2016.
- Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, pp. 2261–2269, 2017.
- Takashi Ishida, Gang Niu, Weihua Hu, and Masashi Sugiyama. Learning from complementary labels. In *NeurIPS*, pp. 5639–5649, 2017.

- Takashi Ishida, Gang Niu, and Masashi Sugiyama. Binary classification from positive-confidence data. In *NeurIPS 2018*, pp. 5921–5932, 2018.
- Takashi Ishida, Gang Niu, Aditya Krishna Menon, and Masashi Sugiyama. Complementary-label learning for arbitrary losses and models. In *ICML*, pp. 2971–2980, 2019.
- Yasuhiro Katsura and Masato Uchida. Bridging ordinary-label learning and complementary-label learning. In *ACML*, pp. 161–176, 2020.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.
- Yu-Feng Li and De-Ming Liang. Safe semi-supervised learning: A brief introduction. *Frontiers Comput. Sci.*, 13(4):669–676, 2019.
- Yu-Feng Li and Zhi-Hua Zhou. Towards making unlabeled data never hurt. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(1):175–188, 2015.
- Nan Lu, Gang Niu, Aditya Krishna Menon, and Masashi Sugiyama. On the minimal supervision for training any binary classifier from only unlabeled data. In *ICLR*, 2019.
- Nan Lu, Tianyi Zhang, Gang Niu, and Masashi Sugiyama. Mitigating overfitting in supervised classification from two unlabeled datasets: A consistent risk correction approach. In *AISTATS*, pp. 1115–1125, 2020.
- Jiaqi Lv, Miao Xu, Lei Feng, Gang Niu, Xin Geng, and Masashi Sugiyama. Progressive identification of true labels for partial-label learning. In *ICML*, pp. 6500–6510, 2020.
- Andreas Maurer. A vector-contraction inequality for rademacher complexities. In Ronald Ortner, Hans Ulrich Simon, and Sandra Zilles (eds.), *ALT*, pp. 3–17, 2016.
- C. McDiarmid. *On the method of bounded differences*. Surveys in Combinatorics, 1989.
- Shahar Mendelson. Lower bounds for the empirical minimization algorithm. *IEEE Trans. Inf. Theory*, 54(8):3797–3803, 2008.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. Adaptive computation and machine learning. MIT Press, 2012. ISBN 978-0-262-01825-8.
- Nagarajan Natarajan, Inderjit S. Dhillon, Pradeep Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *NeurIPS*, pp. 1196–1204, 2013.
- Gang Niu, Wittawat Jitkrittum, Bo Dai, Hirotaka Hachiya, and Masashi Sugiyama. Squared-loss mutual information regularization: A novel information-theoretic approach to semi-supervised learning. In *ICML*, pp. 10–18, 2013.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pp. 8024–8035, 2019.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake VanderPlas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Edouard Duchesnay. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.*, 12:2825–2830, 2011.
- Lukas Ruff, Nico Görnitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Robert A. Vandermeulen, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *ICML*, pp. 4390–4399, 2018.

- Tomoya Sakai, Marthinus Christoffel du Plessis, Gang Niu, and Masashi Sugiyama. Semi-supervised classification based on classification from positive and unlabeled data. In *ICML*, pp. 2998–3006, 2017.
- Emanuele Sansone, Francesco G. B. De Natale, and Zhi-Hua Zhou. Efficient training for positive unlabeled learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(11):2584–2598, 2019.
- Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alexander J. Smola, and Robert C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Comput.*, 13(7):1443–1471, 2001.
- Takuya Shimada, Han Bao, Issei Sato, and Masashi Sugiyama. Classification from pairwise similarities/dissimilarities and unlabeled data via empirical risk minimization. *CoRR*, abs/1904.11717, 2019.
- Kazuhiko Shinoda, Hiroataka Kaji, and Masashi Sugiyama. Binary classification from positive data with skewed confidence. In *IJCAI*, pp. 3328–3334, 2020.
- Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density Ratio Estimation in Machine Learning*. Cambridge University Press, 2012. ISBN 978-0-521-19017-6. URL <http://www.cambridge.org/de/academic/subjects/computer-science/pattern-recognition-and-machine-learning/density-ratio-estimation-machine-learning?format=HB>.
- David M. J. Tax and Robert P. W. Duin. Support vector domain description. *Pattern Recognit. Lett.*, 20(11-13):1191–1199, 1999.
- Ambuj Tewari and Peter L. Bartlett. On the consistency of multiclass classification methods. *J. Mach. Learn. Res.*, 8:1007–1025, 2007. URL <http://dl.acm.org/citation.cfm?id=1390325>.
- Vladimir Vapnik. *Statistical Learning Theory*. Wiley, 1998. ISBN 978-0-471-03003-4.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017. URL <http://arxiv.org/abs/1708.07747>.
- Yixing Xu, Chang Xu, Chao Xu, and Dacheng Tao. Multi-positive and unlabeled learning. In *IJCAI*, pp. 3182–3188, 2017.
- Xiyu Yu, Tongliang Liu, Mingming Gong, and Dacheng Tao. Learning with biased complementary labels. In *ECCV*, pp. 69–85, 2018.
- Mingyuan Zhang, Jane Lee, and Shivani Agarwal. Learning from noisy labels with no change to the training process. In *ICML*, volume 139, pp. 12468–12478, 2021.
- Tong Zhang. Statistical analysis of some multi-category large margin classification methods. *J. Mach. Learn. Res.*, 5:1225–1251, 2004. URL <http://jmlr.org/papers/volume5/zhang04b/zhang04b.pdf>.
- Zhenyu Zhang, Peng Zhao, Yuan Jiang, and Zhi-Hua Zhou. Learning from incomplete and inaccurate supervision. In *KDD*, pp. 1017–1025, 2019.
- Zhi-Hua Zhou. A brief introduction to weakly supervised learning. *National Science Review*, 005(001):44–53, 2018.
- Xiaojin Zhu and Andrew B. Goldberg. *Introduction to Semi-Supervised Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2009.

A APPENDIX

B PROOF OF THEOREM 1

Proof.

$$\begin{aligned}
\pi_{y_s} \mathbb{E}_{p(\mathbf{x}|y_s)} \left[\sum_{y=1}^K \frac{r^y(\mathbf{x})}{r^{y_s}(\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) \right] &= \pi_{y_s} \int \sum_{y=1}^K \frac{p(y|\mathbf{x})}{p(y_s|\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) p(\mathbf{x}|y_s) d\mathbf{x} \\
&= \int \sum_{y=1}^K \frac{p(\mathbf{x}, y_s)}{p(y_s|\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) p(y|\mathbf{x}) d\mathbf{x} \\
&= \int \sum_{y=1}^K p(\mathbf{x}) \ell(\mathbf{g}(\mathbf{x}), y) p(y|\mathbf{x}) d\mathbf{x} \\
&= \int \sum_{y=1}^K \ell(\mathbf{g}(\mathbf{x}), y) p(\mathbf{x}, y) d\mathbf{x} = R(\mathbf{g}).
\end{aligned}$$

□

C PROOF OF THEOREM 2

To prove the Theorem 2, we first give the definition of Rademacher complexity:

Definition 2. (*Rademacher complexity*) Let Z_1, \dots, Z_n be n i.i.d. random variables drawn from a probability distribution μ and $\mathcal{F} = \{f : Z \rightarrow \mathbb{R}\}$ be a class of measurable functions. Then the expected Rademacher complexity of function class \mathcal{F} is given by:

$$\mathfrak{R}_n(\mathcal{F}) = \mathbb{E}_{Z_1, \dots, Z_n \sim \mu} \mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \sigma_i f(Z_i) \right], \quad (14)$$

where $\sigma_1, \dots, \sigma_n$ are the Rademacher variables that take the value from $\{-1, +1\}$ evenly.

Based on the setting and assumptions in Section 3.2, we have the following lemma for bounding the uniform convergence:

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

$$\sup_{\mathbf{g} \in \mathcal{G}} \left| \hat{R}_{SC}(\mathbf{g}) - R(\mathbf{g}) \right| \leq \frac{2\sqrt{2}\pi_{y_s} L_\ell}{C_r} \sum_{y=1}^K \mathfrak{R}_n(\mathcal{G}_y) + \frac{\pi_{y_s} C_\ell}{C_r} \sqrt{\frac{\log \frac{2}{\delta}}{2n}}, \quad (15)$$

where $\mathfrak{R}_n(\mathcal{G}_y)$ is the Rademacher complexity of \mathcal{G}_y on i.i.d. samples with size n drawn from a single class $p(\mathbf{x}|y_s)$.

Proof. We begin with proving that the one direction $\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}_{SC}(\mathbf{g}) - R(\mathbf{g})$ is bounded with probability at least $1 - \frac{\delta}{2}$. Suppose an instance \mathbf{x}_i is changed by \mathbf{x}'_i , we can see that then change of $\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}_{SC}(\mathbf{g}) - R(\mathbf{g})$ is no greater than $\pi_{y_s} C_\ell / n C_r$. By applying the McDiarmid's inequality (McDiarmid, 1989), with probability at least $1 - \frac{\delta}{2}$, the following inequality holds:

$$\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}_{SC}(\mathbf{g}) - R(\mathbf{g}) \leq \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_n} \left[\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}_{SC}(\mathbf{g}) - R(\mathbf{g}) \right] + \frac{\pi_{y_s} C_\ell}{C_r} \sqrt{\frac{\log \frac{2}{\delta}}{2n}}.$$

Denote by $\mathcal{L}(\mathbf{g}(\mathbf{x})) = \sum_{y=1}^K \frac{r^y}{r^{y_s}} * \ell(\mathbf{g}(\mathbf{x}), y)$. It is easy to show that $\mathcal{L}(\mathbf{g}(\mathbf{x}))$ is $\frac{L_\ell}{C_r}$ -Lipschitz w.r.t. $\mathbf{g}(\mathbf{x})$ due to the fact that $\sum_{y=1}^K r^y = 1$ and $r^{y_s} \geq C_r$. Since $\hat{R}_{SC}(\mathbf{g})$ is unbiased, it is routine to

show that (Mohri et al., 2012):

$$\begin{aligned} \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_n} \left[\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}_{SC}(\mathbf{g}) - R(\mathbf{g}) \right] &\leq 2\pi_{y_s} \mathfrak{R}_n(\mathcal{L} \circ \mathcal{G}) \\ &\leq \frac{2\sqrt{2}\pi_{y_s} L_\ell}{C_r} \sum_{y=1}^K \mathfrak{R}_n(\mathcal{G}_y), \end{aligned}$$

The last inequality holds according to the Talagrand’s contraction inequality (Maurer, 2016).

The proof of the other direction $\sup_{\mathbf{g} \in \mathcal{G}} R(\mathbf{g}) - \hat{R}_{SC}(\mathbf{g})$ is similar. Thus we conclude the proof. \square

Then we can begin to prove the Theorem 2:

Proof.

$$\begin{aligned} R(\hat{\mathbf{g}}_{SC}) - R(\mathbf{g}^*) &= \left(R(\hat{\mathbf{g}}_{SC}) - \hat{R}_{SC}(\hat{\mathbf{g}}_{SC}) \right) + \left(\hat{R}_{SC}(\hat{\mathbf{g}}_{SC}) - \hat{R}_{SC}(\mathbf{g}^*) \right) + \left(\hat{R}_{SC}(\mathbf{g}^*) - R(\mathbf{g}^*) \right) \\ &\leq \left(R(\hat{\mathbf{g}}_{SC}) - \hat{R}_{SC}(\hat{\mathbf{g}}_{SC}) \right) + \left(\hat{R}_{SC}(\mathbf{g}^*) - R(\mathbf{g}^*) \right) \\ &\leq 2 \sup_{\mathbf{g} \in \mathcal{G}} \left| R(\hat{\mathbf{g}}) - \hat{R}_{SC}(\hat{\mathbf{g}}) \right|, \end{aligned}$$

where the second inequality holds according to the definition of ERM. Combining this conclusion with the previous lemma, we can conclude the proof. \square

D PROOF OF THEOREM 3

Proof. First of all, we give the definition of classification-calibrated loss:

Proposition 1. (Tewari & Bartlett, 2007) *If the classification-calibrated losses are used, the optimal classifier \mathbf{g}^* that minimizes the classification risk among all the measurable functions is also Bayes-optimal classifier: the classifier that minimizes the classification risk w.r.t. 0-1 loss, i.e., misclassification rate. In other words, the prediction of the classifier on \mathbf{x} satisfies this condition: $\operatorname{argmax}_{y \in \mathcal{Y}} g_y(\mathbf{x}) = \operatorname{argmax}_{y \in \mathcal{Y}} p(y|\mathbf{x})$.*

Denote by $\tilde{p}(\mathbf{x}, y) = \tilde{p}(y|\mathbf{x})p(\mathbf{x})$. Notice that the following holds:

$$\begin{aligned} \int \sum_{y=1}^K \tilde{p}(\mathbf{x}, y) d\mathbf{x} &= \int \sum_{y=1}^K \tilde{p}(y|\mathbf{x})p(\mathbf{x}) d\mathbf{x} \\ &= \int \left(\sum_{y=1}^K \tilde{p}(y|\mathbf{x}) \right) p(\mathbf{x}) d\mathbf{x} \\ &= \int p(\mathbf{x}) d\mathbf{x} = 1 \end{aligned}$$

Then we can see that $\tilde{p}(\mathbf{x}, y)$ is the density of a valid joint distribution on $\mathcal{X} \times \mathcal{Y}$. Then we have:

$$\begin{aligned}\tilde{R}(\mathbf{g}) &= \mathbb{E}_{p(\mathbf{x}|y_s)} \left[\phi(\mathbf{x}) \sum_{y=1}^K \tilde{r}^y(\mathbf{x}) \ell(\mathbf{g}(\mathbf{x}), y) \right] \\ &= \int p(\mathbf{x}|y_s) \phi(\mathbf{x}) \sum_{y=1}^K \tilde{r}^y(\mathbf{x}) \ell(\mathbf{g}(\mathbf{x}), y) d\mathbf{x} \\ &= \int p(\mathbf{x}) \sum_{y=1}^K \tilde{r}^y(\mathbf{x}) \ell(\mathbf{g}(\mathbf{x}), y) d\mathbf{x} \\ &= \int \sum_{y=1}^K \tilde{p}(\mathbf{x}, y) \ell(\mathbf{g}(\mathbf{x}), y) d\mathbf{x} \\ &= \mathbb{E}_{\tilde{p}(\mathbf{x}, y)} [\ell(\mathbf{g}(\mathbf{x}), y)]\end{aligned}$$

Let $\tilde{\mathbf{g}}^* = \operatorname{argmin}_{\mathbf{g} \in \mathcal{G}} \tilde{R}(\mathbf{g})$. According to the definition of the classification-calibrated loss, we can see that $\tilde{\mathbf{g}}^{bayes} = \operatorname{argmin}_{\mathbf{g} \text{ measurable}} \tilde{R}(\mathbf{g})$ is Bayes-optimal. Since $\tilde{\mathbf{g}}^{bayes}$ is in \mathcal{G} , we have that $\tilde{R}(\tilde{\mathbf{g}}^{bayes}) \geq \tilde{R}(\tilde{\mathbf{g}}^*)$. However, since $\tilde{R}(\tilde{\mathbf{g}}^{bayes}) = \min_{\mathbf{g} \text{ measurable}} \tilde{R}(\mathbf{g})$, we have that $\tilde{R}(\tilde{\mathbf{g}}^{bayes}) \leq \tilde{R}(\tilde{\mathbf{g}}^*)$. Combining the two conclusions, we know that $\tilde{R}(\tilde{\mathbf{g}}^{bayes}) = \tilde{R}(\tilde{\mathbf{g}}^*)$. Then we can see that $\tilde{\mathbf{g}}^* = \operatorname{argmin}_{\mathbf{g} \text{ measurable}} \tilde{R}(\mathbf{g})$, which indicates that $\tilde{\mathbf{g}}^*$ is also Bayes-optimal.

According to the definition of the Bayes-optimal classifier, we have that $\operatorname{argmax}_{y \in \mathcal{Y}} \tilde{g}_y^*(\mathbf{x}) = \operatorname{argmax}_{y \in \mathcal{Y}} \tilde{p}(y|\mathbf{x})$. Since $\operatorname{argmax}_{y \in \mathcal{Y}} \tilde{p}(y|\mathbf{x}) = \operatorname{argmax}_{y \in \mathcal{Y}} p(y|\mathbf{x})$, we have that $\operatorname{argmax}_{y \in \mathcal{Y}} \tilde{g}_y^*(\mathbf{x}) = \operatorname{argmax}_{y \in \mathcal{Y}} p(y|\mathbf{x})$. Then we can see that $\tilde{\mathbf{g}}^*$ is also the Bayes-optimal classifier of $R(\mathbf{g})$, which shows the classifier-consistency of $\tilde{R}(\mathbf{g})$. \square

E PROOF OF LEMMA 1 AND THEOREM 4

We begin with the proof of Lemma 1:

Proof. We prove this lemma by bounding the uniform convergence as in the proof of Theorem 2. First we prove the following technical lemma:

Lemma 4. *For any $0 < \delta < 1$, the following inequality holds with probability at least $1 - \delta$:*

$$\sup_{\mathbf{g} \in \mathcal{G}} \left| \hat{R}(\mathbf{g}) - \tilde{R}(\mathbf{g}) \right| \leq \frac{2\sqrt{2}L_\ell}{C_r} \sum_{y=1}^K \mathfrak{R}_n(\mathcal{G}_y) + \frac{C_\ell}{C_r} \sqrt{\frac{\log \frac{2}{\delta}}{2n}}, \quad (16)$$

where $\mathfrak{R}_n(\mathcal{G}_y)$ is the Rademacher complexity of \mathcal{G}_y on i.i.d. samples with size n drawn from a single class $p(\mathbf{x}|y_s)$.

Proof. We only prove that the direction $\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}(\mathbf{g}) - \tilde{R}(\mathbf{g})$ is bounded with probability at least $1 - \frac{\delta}{2}$ since the proof of another direction is completely symmetric. Suppose an instance \mathbf{x}_i is changed by \mathbf{x}'_i , we can see that then change of $\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}(\mathbf{g}) - \tilde{R}(\mathbf{g})$ is no greater than C_ℓ/nC_r since $\phi_i \leq 1/C_r$. By applying the McDiarmid's inequality (McDiarmid, 1989), with probability at least $1 - \frac{\delta}{2}$, the following inequality holds:

$$\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}(\mathbf{g}) - \tilde{R}(\mathbf{g}) \leq \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_n} \left[\sup_{\mathbf{g} \in \mathcal{G}} \hat{R}(\mathbf{g}) - \tilde{R}(\mathbf{g}) \right] + \frac{C_\ell}{C_r} \sqrt{\frac{\log \frac{2}{\delta}}{2n}}.$$

Denote by $\tilde{\mathcal{L}}(\mathbf{g}(\mathbf{x})) = \phi(\mathbf{x}) \sum_{y=1}^K \tilde{r}^y * \ell(\mathbf{g}(\mathbf{x}), y)$. It is easy to show that $\mathcal{L}(\mathbf{g}(\mathbf{x}))$ is $\frac{L_\ell}{C_r}$ -Lipschitz w.r.t. $\mathbf{g}(\mathbf{x})$ due to the fact that $\sum_{y=1}^K r^y = 1$ and $\phi(\mathbf{x}) \leq 1/C_r$. Then it is routine to show that

Since the factor $\inf_{\mathbf{x} \in \mathcal{X}} \Delta(\mathbf{x})$ is irrelevant to \mathbf{x} , we can further give the following conclusion:

$$\mathbb{E}_{p(\mathbf{x})} [1(f(\mathbf{x}) \neq \operatorname{argmax}_{y \in \mathcal{Y}} \tilde{p}(y|\mathbf{x})))] \leq \frac{1}{\inf_{\mathbf{x} \in \mathcal{X}} \Delta(\mathbf{x})} \left(\tilde{R}_{01}(\hat{\mathbf{g}}) - \tilde{R}_{01}(\tilde{\mathbf{g}}^*) \right). \quad (17)$$

Since $\tilde{\mathbf{g}}^* \in \mathcal{G}$ is also the Bayes-optimal classifier of $R(\mathbf{g})$ as stated in the proof of Theorem 3, we can learn that $R_{01}(\tilde{\mathbf{g}}^*) \leq R_{01}(\mathbf{g}_{01}^*)$. Notice that $\tilde{\mathbf{g}}^*$ is in \mathcal{G} , then $R_{01}(\tilde{\mathbf{g}}^*) \geq R_{01}(\mathbf{g}_{01}^*)$. Combining the two inequalities, we can see that $R_{01}(\tilde{\mathbf{g}}^*) = R_{01}(\mathbf{g}_{01}^*)$. Then we can get the expression of $R_{01}(\hat{\mathbf{g}}) - R_{01}(\mathbf{g}_{01}^*)$:

$$\begin{aligned} R_{01}(\hat{\mathbf{g}}) - R_{01}(\mathbf{g}_{01}^*) &= R_{01}(\hat{\mathbf{g}}) - R_{01}(\tilde{\mathbf{g}}^*) \\ &= \mathbb{E}_{p(\mathbf{x}, y)} [1(f(\mathbf{x}) \neq y)] - \mathbb{E}_{p(\mathbf{x}, y)} [1(\operatorname{argmax}_{y \in \mathcal{Y}} \tilde{g}_y^* \neq y)] \\ &= \mathbb{E}_{p(\mathbf{x})} [\mathbb{E}_{p(y|\mathbf{x})} [1(f(\mathbf{x}) \neq y) - 1(\operatorname{argmax}_{y \in \mathcal{Y}} \tilde{g}_y^* \neq y)]] \\ &= \mathbb{E}_{p(\mathbf{x})} \left[1 - p(f(\mathbf{x})|\mathbf{x}) - \left(1 - \max_{y \in \mathcal{Y}} p(y|\mathbf{x}) \right) \right] \\ &= \mathbb{E}_{p(\mathbf{x})} \left[\max_{y \in \mathcal{Y}} p(y|\mathbf{x}) - p(f(\mathbf{x})|\mathbf{x}) \right] \\ &= \mathbb{E}_{p(\mathbf{x})} \left[\max_{y \in \mathcal{Y}} p(y|\mathbf{x}) - p(f(\mathbf{x})|\mathbf{x}) \right] \\ &= \mathbb{E}_{p(\mathbf{x})} \left[1(f(\mathbf{x}) \neq \operatorname{argmax}_{y \in \mathcal{Y}} p(y|\mathbf{x})) \left(\max_{y \in \mathcal{Y}} p(y|\mathbf{x}) - p(f(\mathbf{x})|\mathbf{x}) \right) \right]. \end{aligned}$$

Since $\operatorname{argmax}_{y \in \mathcal{Y}} p(y|\mathbf{x}) = \operatorname{argmax}_{y \in \mathcal{Y}} \tilde{p}(y|\mathbf{x})$ and $(\max_{y \in \mathcal{Y}} p(y|\mathbf{x}) - p(f(\mathbf{x})|\mathbf{x})) \leq 1$, we can get the following inequality:

$$R_{01}(\hat{\mathbf{g}}) - R_{01}(\mathbf{g}_{01}^*) \leq \mathbb{E}_{p(\mathbf{x})} [1(f(\mathbf{x}) \neq \operatorname{argmax}_{y \in \mathcal{Y}} \tilde{p}(y|\mathbf{x}))].$$

Combining the inequality above, (17), and Lemma 2, we can conclude the proof of Theorem 4. \square

F PROOF OF THEOREM 5

Proof. $\hat{B}_\eta^{SC}(\hat{\phi}) = \frac{1}{n} \sum_{i=1}^n \left(\nabla \eta(\hat{\phi}(\mathbf{x}_i)) \hat{\phi}(\mathbf{x}_i) - \eta(\hat{\phi}(\mathbf{x}_i)) \right)$ and $\hat{B}_\eta^U(\hat{\phi}) = -\frac{1}{n_u} \sum_{i=1}^{n_u} \nabla \eta(\hat{\phi}(\mathbf{x}_i^u))$. $\hat{B}_\eta^{SC}(\hat{\phi})$ and $\hat{B}_\eta^U(\hat{\phi})$ are their expectations, respectively. Then the following work is to bound the uniform convergence. Assume that there is $C_{\hat{\phi}} > 0$ that $\sup_{\hat{\phi} \in \Phi} \|\hat{\phi}\|_\infty \leq C_{\hat{\phi}}$. The function $\eta(\hat{\phi}(\mathbf{x}))$ is Lipschitz continuous for all $\|\hat{\phi}\|_\infty \leq C_{\hat{\phi}}$ with Lipschitz constant $L_\eta > 0$ and upper-bounded by $C_\eta > 0$. $\nabla \eta$ is also Lipschitz continuous with constant $L'_\eta > 0$. Then we first give the technical lemma:

Lemma 5. *For any $0 < \delta < 1$, With probability at least $1 - \delta$:*

$$\sup_{\hat{\phi} \in \Phi} |\hat{B}_\eta^{SC}(\hat{\phi}) - B_\eta^{SC}(\hat{\phi})| \leq 2L_B \mathfrak{R}_n(\Phi) + C_B \sqrt{\frac{\log \frac{2}{\delta}}{2n}}. \quad (18)$$

$$\sup_{\hat{\phi} \in \Phi} |\hat{B}_\eta^U(\hat{\phi}) - B_\eta^U(\hat{\phi})| \leq 2L'_\eta \mathfrak{R}_{n_u}^u(\Phi) + L_\eta \sqrt{\frac{\log \frac{2}{\delta}}{2n_u}}. \quad (19)$$

Proof. Notice that $\hat{B}_\eta^{SC}(\hat{\phi})$ is Lipschitz continuous with Lipschitz constant $L_B = L'_\eta C_{\hat{\phi}} + 2L_\eta$ and its absolute value is bounded by $C_B = \max\{C_\eta, L_\eta C_{\hat{\phi}}\}$. By using the McDiarmid's inequality, Union bound, and the Talagrand's contraction inequality, we can learn that with probability at least $1 - \delta$:

$$\sup_{\hat{\phi} \in \Phi} |\hat{B}_\eta^{SC}(\hat{\phi}) - B_\eta^{SC}(\hat{\phi})| \leq 2L_B \mathfrak{R}_n(\Phi) + C_B \sqrt{\frac{\log \frac{2}{\delta}}{2n}}.$$

Notice that $\hat{B}_\eta^U(\hat{\phi})$ is L'_η -Lipschitz continuous and is bounded by L_η . In the same way, we can see that with probability at least $1-\delta$:

$$\sup_{\hat{\phi} \in \Phi} |\hat{B}_\eta^U(\hat{\phi}) - B_\eta^U(\hat{\phi})| \leq 2L'_\eta \mathfrak{R}_{n_u}^u(\Phi) + L_\eta \sqrt{\frac{\log \frac{2}{\delta}}{2n_u}}.$$

□

Then we can begin to prove the Theorem 5:

$$\begin{aligned} B_\eta(\phi \|\hat{\phi}^*) &= B_\eta(\phi \|\hat{\phi}^*) - B_\eta(\phi \|\phi) \\ &= B_\eta(\phi \|\hat{\phi}^*) - \hat{B}_\eta(\phi \|\hat{\phi}^*) + \hat{B}_\eta(\phi \|\hat{\phi}^*) - \hat{B}_\eta(\phi \|\phi) + \hat{B}_\eta(\phi \|\phi) - B_\eta(\phi \|\phi) \\ &\leq 2 \sup_{\hat{\phi} \in \Phi} |B_\eta(\phi \|\hat{\phi}) - \hat{B}_\eta(\phi \|\hat{\phi})| \\ &\leq 2 \sup_{\hat{\phi} \in \Phi} \left(|\hat{B}_\eta^{SC}(\hat{\phi}) - B_\eta^{SC}(\hat{\phi})| + |\hat{B}_\eta^U(\hat{\phi}) - B_\eta^U(\hat{\phi})| \right) \\ &\leq 2 \sup_{\hat{\phi} \in \Phi} |\hat{B}_\eta^{SC}(\hat{\phi}) - B_\eta^{SC}(\hat{\phi})| + 2 \sup_{\hat{\phi} \in \Phi} |\hat{B}_\eta^U(\hat{\phi}) - B_\eta^U(\hat{\phi})|. \end{aligned}$$

According to the technical lemma above and the union bound, we can show that with probability at least $1 - \delta$:

$$B_\eta(\phi \|\hat{\phi}^*) \leq 4L_B \mathfrak{R}_n(\Phi) + 2C_B \sqrt{\frac{\log \frac{4}{\delta}}{2n}} + 4L'_\eta \mathfrak{R}_{n_u}^u(\Phi) + 2L_\eta \sqrt{\frac{\log \frac{4}{\delta}}{2n_u}}.$$

which concludes the proof. □

G PROOF OF THEOREM 6 AND THE FORMULATION OF NOISE-ROBUST SUB-CONF LEARNING

Proof.

$$\begin{aligned} &\left(\sum_{y \in \mathcal{Y}_s} \pi_y \right) \mathbb{E}_{p(\mathbf{x}|y \in \mathcal{Y}_s)} \left[\sum_{y=1}^K \left(\frac{p(y|\mathbf{x})}{\sum_{y_s \in \mathcal{Y}_s} p(y_s|\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) \right) \right] \\ &= \left(\sum_{y \in \mathcal{Y}_s} \pi_y \right) \int \sum_{y=1}^K \left(\frac{p(y|\mathbf{x})}{p(y \in \mathcal{Y}_s|\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) \right) p(\mathbf{x}|y \in \mathcal{Y}_s) d\mathbf{x} \\ &= \left(\sum_{y \in \mathcal{Y}_s} \pi_y \right) \int \sum_{y=1}^K \left(\frac{p(\mathbf{x}|y \in \mathcal{Y}_s)}{p(y \in \mathcal{Y}_s|\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) \right) p(y|\mathbf{x}) d\mathbf{x} \\ &= \int \sum_{y=1}^K \left(\frac{p(\mathbf{x}, y \in \mathcal{Y}_s)}{p(y \in \mathcal{Y}_s|\mathbf{x})} \ell(\mathbf{g}(\mathbf{x}), y) \right) p(y|\mathbf{x}) d\mathbf{x} \\ &= \int \sum_{y=1}^K (p(\mathbf{x}) \ell(\mathbf{g}(\mathbf{x}), y)) p(y|\mathbf{x}) d\mathbf{x} \\ &= \int \sum_{y=1}^K \ell(\mathbf{g}(\mathbf{x}), y) p(\mathbf{x}) p(y|\mathbf{x}) d\mathbf{x} \\ &= R(\mathbf{g}). \end{aligned}$$

□

Then we begin to give the formulation of Noise-Robust Sub-Conf learning:

Definition 3. *The risk estimator of Noise-Robust Sub-Conf learning is defined as follow:*

$$\hat{R}_{sub}(\mathbf{g}) = \frac{1}{n} \sum_{i=1}^n \phi_i \sum_{y=1}^K \tilde{r}_i^y \ell(\mathbf{g}(\mathbf{x}_i), y). \quad (20)$$

where $\phi(\mathbf{x}) = \frac{p(\mathbf{x})}{p(\mathbf{x}|y \in \mathcal{Y}_s)}$ and $\phi_i = \frac{p(\mathbf{x}_i)}{p(\mathbf{x}_i|y \in \mathcal{Y}_s)}$.

Notice that if we substitute $p(\mathbf{x}|y \in \mathcal{Y}_s)$ with $p(\mathbf{x}|y_s)$, (20) is converted to (6). We can also get the infinite and finite-sample consistency of learning with (20) by conducting such substitution.

H ADDITIONAL INFORMATION OF BENCHMARK EXPERIMENTS

H.1 DETAILED INFORMATION OF BENCHMARK DATASETS

In Section 6, we used 2 widely-used large-scale benchmark datasets. Here, we report the sources of these datasets and the way we split them.

- Fashion-MNIST (Xiao et al., 2017). It is a 10-class dataset of fashion items. Each instance is a 28*28 grayscale image. Source: <https://github.com/zalandoresearch/fashion-mnist>.
- CIFAR-10 (Krizhevsky et al., 2009). It is a 10-class dataset for 10 different objects and each instance is a 32*32*3 colored image in RGB format. Source: <https://www.cs.toronto.edu/~kriz/cifar.html>.

In the experiments, we first split 30% of the dataset for generating confidences. Then we further split 10% of the dataset for density ratio estimation. Finally we use 10% of the dataset as the validation set. The rest of the dataset are used as the training set.

H.2 DETAILED INFORMATION OF THE MODELS AND OPTIMIZATION ALGORITHM

For generating confidence scores, the model used for Fashion-MNIST is a 3-layer MLP (d-100-100-100-10) with ReLU and Resnet-18 (He et al., 2016) is used for CIFAR-10. For density ratio estimation, the output dimension is changed to 1 and other settings are the same. For classification, a 3-layer MLP (d-500-500-500-10) with ReLU is used for Fashion-MNIST, and DenseNet-161 is used for CIFAR-10.

Adam with default momentum was used for optimization in this paper. For generating confidence scores, the epoch number, batch size, learning rate, and weight decay are set to 20, 100, 1e-4, and 1e-4 for Fashion-MNIST, respectively. For density ratio estimation, the epoch number is reduced to 10 and other hyperparameters are the same. When training the classifiers, the epoch number is increased to 100. For CIFAR-10, the learning rate and weight decay are changed to 1e-3 and other hyperparameters are the same as those of Fashion-MNIST.

I EXPERIMENTAL RESULTS ON MORE BENCHMARK DATASETS

In this section, we show the performance of our methods and baseline methods on more benchmark datasets. We use two widely-used benchmark datasets and report the sources of these datasets. the way we split them is the same as that in the previous section.

- MNIST (LeCun et al., 1998). It is a grayscale dataset of handwritten digits from 0 to 9, where the size of the images is 28*28. Source: <http://yann.lecun.com/exdb/mnist/>.
- Kuzushiji-MNIST (Clanuwat et al., 2018). It is a 10-class dataset of cursive Japanese characters ('Kuzushiji'). Source: <https://github.com/rois-codh/kmnist>.

For generating confidence scores, the model used is a 3-layer MLP (d-100-100-100-10) with ReLU. For density ratio estimation, the output dimension is changed to 1 and other settings are the same. For classification, a 3-layer MLP (d-500-500-500-10) with ReLU is used.

Table 3: Mean and standard deviation of the classification accuracy over 20 trials for the MNIST dataset. The proposed methods were compared with the baseline Weighted method and fully-supervised method, with different classes used for training. The best and equivalent methods are shown in bold based on the 5% t-test, excluding fully-supervised method.

Used Classes		SC/Sub-Conf	NoRSC-Conf	Weighted	Supervised
8	Accurate	89.52±0.60	87.87±1.27	87.68±1.20	94.82±0.53
	Noisy	-±-	68.33±3.62	63.39±2.65	
7 & 9	Accurate	88.46±0.92	80.08±1.86	84.13±2.00	
	Noisy	-±-	73.48±1.43	71.61±1.78	
0 & 2 & 3	Accurate	91.20±0.26	88.43±0.53	90.66±0.18	
	Noisy	-±-	88.33±0.15	88.32±0.47	
1 & 5 & 7 & 9	Accurate	91.28±0.21	90.89±0.16	89.72±0.45	
	Noisy	-±-	86.50±0.72	83.85±1.20	

Adam with default momentum was used for optimization in this paper. For generating confidence scores, the epoch number, batch size, learning rate, and weight decay are set to 20, 1000, 1e-4, and 1e-4, respectively. For density ratio estimation, the epoch number is reduced to 10 and other hyperparameters are the same. When training the classifiers, the epoch number is increased to 100. The experimental results are shown in Table 3 and 4.

It can be seen from the experimental results that our SC/Sub-Conf learning method outperform other baseline methods in most cases and the NoRSC-Conf learning method can remain effective under extreme noise, which supports our claims in Section 6.

Table 4: Mean and standard deviation of the classification accuracy over 20 trials for the Kuzushiji-MNIST dataset. The proposed methods were compared with the baseline Weighted method and fully-supervised method, with different classes used for training. The best and equivalent methods are shown in bold based on the 5% t-test, excluding fully-supervised method.

Used Classes		SC/Sub-Conf	NoRSC-Conf	Weighted	Supervised
'Tsu'	Accurate	68.45±1.42	67.33±1.08	64.55±1.28	76.06±0.32
	Noisy	-±-	38.14±2.12	34.99±2.27	
'Ha'	Accurate	68.19±0.56	67.07±0.65	66.86±0.67	
	Noisy	-±-	51.67±1.26	48.21±1.63	
'Ma'	Accurate	66.20±0.93	63.67±1.17	63.56±0.96	
	Noisy	-±-	38.84±2.41	35.80±2.35	
'Su' & 'Na'	Accurate	69.98±0.72	67.85±0.40	68.99±0.59	
	Noisy	-±-	62.17±1.26	61.95±1.11	
'Ki' & 'Ya' & 'Re'	Accurate	70.23±0.59	68.64±0.62	69.77±1.58	
	Noisy	-±-	64.60±1.09	64.47±0.88	
'Ki' & 'Su' & 'Ha' & 'Re'	Accurate	70.48±0.47	70.08±0.51	68.56±0.47	
	Noisy	-±-	66.86±0.74	65.42±0.81	